Frameworks with forced symmetry II: Orientation-preserving crystallographic groups

Justin Malestein[*] Louis Theran[†]

Abstract

We give a combinatorial characterization of minimally rigid planar frameworks with orientation-preserving crystallographic symmetry, under the constraint of forced symmetry. The main theorems are proved by extending the methods of the first paper in this sequence from groups generated by a single rotation to groups generated by translations and rotations. The proofs make use of a new family of matroids defined on crystallographic groups and associated submodular functions.

1. Introduction

A *crystallographic framework* is an *infinite* planar structure, *symmetric* with respect to a crystallographic group, made of fixed-length bars connected by universal joints with full rotational freedom. The allowed continuous motions preserve the lengths and connectivity of the bars (as in the finite framework case) and *symmetry with respect to the group* Γ (this is the new addition). However, the representation of Γ is *not* fixed and may change. This model extends the one from [4], using a formalism similar to [13–16]. Figures 1 and 2 show examples of crystallographic frameworks. A crystallographic framework is *rigid* when the only allowed motions (that, additionally, must act on the representation of Γ) are Euclidean isometries and *flexible* otherwise.

1.1 Algebraic setup and combinatorial model

A Γ-crystallographic framework is given by the data $(\tilde{G}, \varphi, \tilde{\boldsymbol{\ell}})$. The infinite graph $\tilde{G}$ encodes the combinatorial structure of the bars. The crystallographic group Γ, along with the free Γ-action φ on $\tilde{G}$ by automorphisms (i.e., $\varphi : \Gamma \to \mathrm{Aut}(\tilde{G})$ is a representation) determines the framework's symmetry; for convenience, we define the notation $\gamma(i) := \varphi(\gamma)(i)$ for $\gamma \in \Gamma$ and $i \in \tilde{V}$. The rest of the framework's geometric data is given by the vector $\tilde{\boldsymbol{\ell}}$, which is an assignment of a positive length to each edge $ij \in \tilde{E}$.

We assume that $\tilde{G}$ has finite quotient $G = \tilde{G}/\Gamma$ with n vertices and m edges. To keep the terminology in this framework manageable, we will refer simply to *frameworks* when the context is clear, with the understanding that the frameworks appearing in the paper are crystallographic.

A *realization* $G(\mathbf{p}, \Phi)$ of the abstract framework $(\tilde{G}, \varphi, \tilde{\boldsymbol{\ell}})$ is defined to be an assignment $\mathbf{p} = (\mathbf{p}_i)_{i \in \tilde{V}}$ of points to the vertices of $\tilde{G}$ and a representation Φ from Γ to a Euclidean isometry

[*]Einstein Institute of Mathematics, Hebrew University of Jerusalem, justinmalestein@gmail.com
[†]Institut für Mathematik, Diskrete Geometrie, Freie Universität Berlin, theran@math.fu-berlin.de

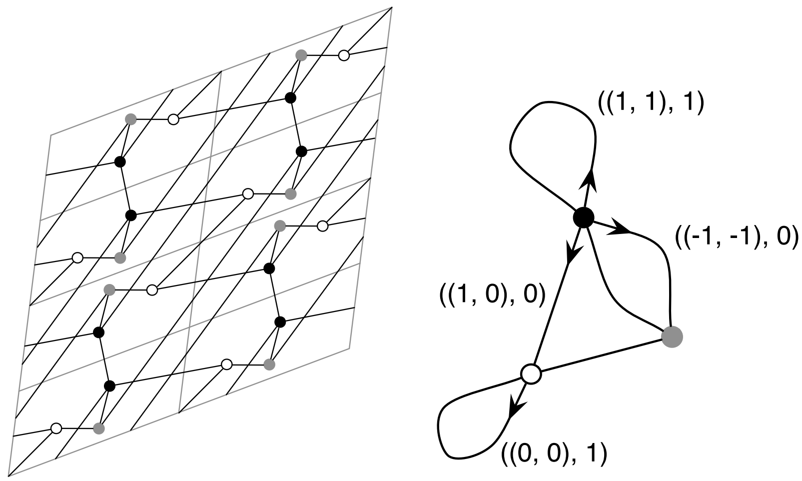

(a) (b)

Figure 1: A Γ_2-crystallographic framework: (a) A piece of an infinite crystallographic framework with Γ_2 symmetry. The group Γ_2 is generated by an order 2 rotation and translations. The origin, which is a rotation center, is at the center of the diagram. Each quadrilateral (with gray edges) is a fundamental domain of the Γ_2-action on $\mathbb{R}^2$. (b) The associated *colored graph* capturing the underlying combinatorics. Edges that are not marked and oriented are colored with the identity element of Γ_2. The vertices in (b) are shaded differently to show the fibers over each of them in (a).

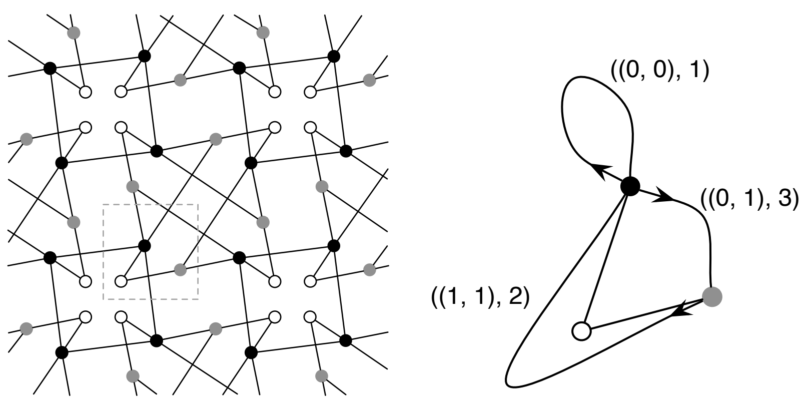

(a) (b)

Figure 2: A Γ_4-crystallographic framework: (a) A piece of an infinite crystallographic framework with Γ_4 symmetry. The group Γ_4 is generated by an order 4 rotation and translations. The fundamental domain of the Γ_4-action on $\mathbb{R}^2$ is shown as a dashed box. (b) The associated colored graph capturing the underlying combinatorics. The color coding conventions are as in Figure 1.

group, such that

$$||\mathbf{p}_i - \mathbf{p}_j|| = \tilde{\ell}_{ij} \qquad \text{for all edges } ij \in \tilde{E} \tag{1}$$

$$\Phi(\gamma) \cdot \mathbf{p}_i = \mathbf{p}_{\gamma(i)} \qquad \text{for all group elements } \gamma \in \Gamma \text{ and vertices } i \in \tilde{V} \tag{2}$$

The condition (1), which appears in the theory of finite frameworks, says that a realization respects the given edge lengths. Equation (2) says that, if we hold Φ fixed, regarded as a map $\mathbf{p} : \tilde{V} \to \mathbb{R}^2$, $\mathbf{p}$ is equivariant with respect to the Γ-actions on $\tilde{G}$ and $\mathbb{R}^2$. However, Φ is, in general, *not* fixed. This is a very important feature of the model: the motions available to the framework include those that deform the representation Φ of Γ, provided this happens in a way compatible with the Γ-action φ.

The *realization space* $\mathcal{R}(\tilde{G}, \varphi, \tilde{\boldsymbol{\ell}})$ (shortly $\mathcal{R}$) of an abstract framework is defined as the set of its realizations. The *configuration space* $\mathcal{C}$ is defined to be the quotient of $\mathcal{R}$ by Euclidean isometries. A realization $\tilde{G}(\mathbf{p}, \Phi)$ is *rigid* if it is isolated in $\mathcal{C}$ and otherwise *flexible*. (See Section 5.1 for a detailed treatment of these spaces.)

As the combinatorial model for crystallographic frameworks it will be more convenient to use colored graphs. A *colored graph* $(G, \boldsymbol{\gamma})$ is a finite, directed graph G, with an assignment $\boldsymbol{\gamma} = (\gamma_{ij})_{ij \in E(G)}$ of an element of a group Γ to each edge.

A straightforward specialization of covering space theory, described in Section 3.1, associates $(\tilde{G}, \varphi)$ with a colored graph $(G, \boldsymbol{\gamma})$: G is the quotient of $\tilde{G}$ by Γ, and the colors encode the covering map $\tilde{G} \to G$ via a map $\rho : \pi_1(G, b) \to \Gamma$.

1.2 Main theorem

Our main result is the following "Maxwell-Laman-type" theorem for crystallographic frameworks where the symmetry group is generated by translations and a finite order rotation. The "Γ-*colored-Laman graphs*" appearing in the statement are defined in Section 3.4; genericity is defined in detail in Section 5.2, but the term is used in the standard sense of algebraic geometry: generic frameworks are the (open, dense) complement of a proper algebraic subset of $\mathbb{R}^m$.

Theorem 1. *Let Γ be an orientation-preserving crystallographic group. A generic crystallographic framework $(\tilde{G}, \varphi, \tilde{\boldsymbol{\ell}})$ with symmetry group Γ is minimally rigid if and only if its colored quotient graph is Γ-colored-Laman.*

Whether a colored graph is Γ-colored-Laman can be checked in polynomial time by combinatorial algorithms based on Edmonds's augmenting path algorithm for Matroid Union [7].

1.3 Infinitesimal rigidity and direction networks

In order to prove the rigidity Theorem 1, we will prove a combinatorial characterization of *generic infinitesimal rigidity*, which is a linearization of the problem. Standard kinds of arguments, along the lines of [1], imply that, generically, rigidity and infinitesimal rigidity coincide. We will study infinitesimal rigidity using crystallographic direction networks.

A Γ-*crystallographic direction network* $(\tilde{G}, \varphi, \tilde{\mathbf{d}})$ consists of an infinite graph $\tilde{G}$ with a free Γ-action φ on the edges and vertices, and an assignment of a *direction* $\tilde{\mathbf{d}}_{ij} \in \mathbb{R}^2 \setminus \{0\}$ to each edge $ij \in \tilde{E}$. We define a realization $G(\mathbf{p}, \Phi)$ of $(\tilde{G}, \varphi, \tilde{\mathbf{d}})$ to be a mapping of $\tilde{V}$ to a point set $\mathbf{p}$ and a

representation Φ of Γ by Euclidean isometries such that

$$\left\langle \mathbf{p}_i - \mathbf{p}_j, \tilde{\mathbf{d}}_{ij}^{\perp} \right\rangle = 0 \qquad \text{for all edges } ij \in \tilde{E} \tag{3}$$

$$\Phi(\gamma) \cdot \mathbf{p}_i = \mathbf{p}_{\gamma(i)} \qquad \text{for all group elements } \gamma \in \Gamma \text{ and vertices } i \in \tilde{V} \tag{4}$$

Equation (3) says that, in any realization, $\mathbf{p}_i - \mathbf{p}_j$ is a scalar multiple of $\tilde{\mathbf{d}}_{ij}$, for each edge $ij \in \tilde{E}$; (4) gives the symmetry constraint. Since setting all the $\mathbf{p}_i$ equal and Φ to be trivial produces a realization, the realization space is never empty. For our purpose, though, such realizations are degenerate. We define a realization of a crystallographic direction network to be *faithful* if none of the edges of $\tilde{G}$ are realized with coincident endpoints. Our second main result is an exact characterization of when a generic direction network admits a faithful realization.

Theorem 2. *Let Γ be an orientation-preserving crystallographic group. A generic Γ-crystallographic direction network $(\tilde{G}, \varphi, \tilde{\mathbf{d}})$ has a unique, up to translation and scaling, faithful realization if and only if its associated colored graph is Γ-colored-Laman.*

1.4 Roadmap and novelty

The overall methodology is an adaptation of the *direction network method* (cf. [24] and [26, Section 4]) for proving rigidity characterizations in the plane. The reduction from infinitesimal rigidity to direction network realizability is, by now, fairly standard. Thus, most of the novelty lies in proving Theorem 2. This is done in three main steps: (i) the construction of a matroid on an orientation-preserving crystallographic group (Section 2); (ii) an extension of the group matroid to one on graphs that serves as a kind of "generalized graphic matroid" (Section 3); (iii) a linear representation result relating bases of the new combinatorial matroid to direction networks with only trivial "collapsed" realizations (Section 4). The last step uses a new kind of geometric argument that is not a straightforward reduction to the Matroid Union Theorem as in [24]; the matroids constructed here, to our knowledge, appear for the first time here (and in [13]).

1.5 History and related work

This paper is the second in a sequence derived from the preprints [13, 14], and the material here has appeared, with the same proofs in [13]. The first part is the submitted manuscript [15]. The results here are built on the theory we developed for studying *periodic frameworks* in [16], which contains a detailed discussion of motivations and other work on periodic frameworks.

The general area of rigidity with symmetry has been somewhat active in the past few years. For completeness, we review some work along similar lines. A specialization of our [16, Theorem A] is due to Ross [19]. Schulze [21, 22] and Schulze and Whiteley [23] studied the question of when "incidental" symmetry induces non-generic behaviors in finite frameworks, which is a different setting than the forced symmetry we consider here and in [15, 16], however one can interpret some of those results in the present setting. Ross, Schulze, and Whiteley [20] have studied the problem we do here, but they do not give any combinatorial characterizations. Borcea and Streinu [5] have proposed a kind of "doubly generic" periodic rigidity, where the combinatorial model does not include the colors on the quotient graph.

A recent preprint of Tanigawa [25] proves a number of parallel redrawing and body-bar rigidity characterizations in higher dimensions and for a larger number of groups than considered

here. The method of [25] is, essentially, to axiomatize the properties of the rank function of the matroid we construct in Section 2.6 and then follow a similar program, making use of a new generalization of Matroid Union.

1.6 Acknowledgements

We thank Igor Rivin for encouraging us to take on this project and many productive discussions on the topic. Our initial work on this topic was part of a larger effort to understand the rigidity and flexibility of hypothetical zeolites, which is supported by NSF CDI-I grant DMR 0835586 to Rivin and M. M. J. Treacy. LT is funded by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no 247029-SDModels. JM is supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no 226135.

2. Groups and matroids

2.1 Crystallographic group preliminaries

We first review some basic facts about orientation-preserving crystallographic groups.

2.1.1. Facts about the Euclidean group The Euclidean isometry group $\mathrm{Euc}(d)$ in any dimension can be represented as the semidirect product $\mathbb{R}^d \rtimes O(d)$ where $O(d)$ is the orthogonal group acting on $\mathbb{R}^d$ in the standard way. The group operation is thus:

$$(\mathbf{v}, r) \cdot (\mathbf{v}', r') = (\mathbf{v} + r \cdot \mathbf{v}', rr')$$

The subgroup $\mathbb{R}^d < \mathrm{Euc}(d)$ is the translation subgroup, and the projection $\mathrm{Euc}(d) \to O(d)$ is the map that associates to an isometry ψ its derivative at the origin $D\psi_0$. The action of $(\mathbf{v}, r) \in \mathrm{Euc}(d)$ on a point $\mathbf{p} \in \mathbb{R}^d$ is $(\mathbf{v}, r) \cdot \mathbf{p} = \mathbf{v} + r \cdot \mathbf{p}$.

Since our setting is 2-dimensional, from now on, we are interested in Euc(2). In the two dimensional case, we have the following simple lemma, which we state without proof.

Lemma 2.1. *Any nontrivial orientation-preserving isometry of the Euclidean plane is either a rotation around a point or a translation.*

Thus, when we refer to orientation-preserving elements of Euc(2) we call them simply *"rotations"* or *"translations"*. We denote the counterclockwise rotation around the origin through angle $2\pi/k$ by R_k.

2.1.2. Crystallographic groups A 2-*dimensional crystallographic group* Γ is a group admitting a discrete cocompact faithful representation $\Gamma \to \mathrm{Euc}(2)$. We will denote by Φ such representations of Γ. In this paper, we are interested in the case where all the group elements are represented by rotations and translations (i.e., we disallow reflections and glides).

The enumeration of the 2-dimensional crystallographic groups is classical, and there are precisely five orientation-preserving ones (see, e.g., [6]). The first group, which we denote by Γ_1, is $\mathbb{Z}^2$. The rest are all semidirect products of $\mathbb{Z}^2$ with a cyclic group. Namely, for $k = 2, 3, 4, 6$, we have $\Gamma_k = \mathbb{Z}^2 \rtimes \mathbb{Z}/k\mathbb{Z}$. The action on $\mathbb{Z}^2$ by the generator of $\mathbb{Z}/k\mathbb{Z}$ is given by the following table.

k	2	3	4	6
matrix	$\begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix}$	$\begin{pmatrix} 0 & -1 \\ 1 & -1 \end{pmatrix}$	$\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$	$\begin{pmatrix} 0 & -1 \\ 1 & 1 \end{pmatrix}$

We define the $\mathbb{Z}^2$ subgroup of Γ_k to be the *translation subgroup* of Γ_k and denote it by $\Lambda(\Gamma_k)$. We denote $\gamma \in \Gamma_k$, $k = 2, 3, 4, 6$ as $\gamma = (t, r)$ with $t \in \mathbb{Z}^2$ and $r \in \mathbb{Z}/k\mathbb{Z}$.

2.1.3. Remark on groups considered Since we are only interested in crystallographic groups of this form, the rest of the paper will consider Γ_k only (and not more general crystallographic groups). Moreover, we will treat only $k = 2, 3, 4, 6$ in what follows, because the case of Γ_1 is covered by [16, Theorem A]. However, the theory and proof methods presented here specialize to Γ_1.

2.1.4. Finitely generated subgroups If $\gamma_1, \dots, \gamma_t$ are element of Γ_k, we denote the subgroup generated by the γ_i as $\langle \gamma_1, \dots, \gamma_t \rangle$. If $\Gamma^1, \dots, \Gamma^t$ are a sequence of finitely generated subgroups then $\langle \Gamma^1, \Gamma^2, \dots, \Gamma^t \rangle$ denotes the subgroup generated by the union of some choice of generators for each Γ^i. We will sometimes abuse notation and consider groups generated together by some elements and some subgroups, e.g. $\langle \gamma_1, \gamma_2, \Gamma^1, \Gamma^2, \Gamma^3 \rangle$.

2.2 Representation space

Γ-crystallographic frameworks and direction networks are required to be symmetric with respect to the group Γ. However, the representation is allowed to flex. In this section, we formalize this flexing.

2.2.1. The representation space Let Γ be a crystallographic group. We define the *representation space* $\mathrm{Rep}(\Gamma)$ of Γ to be

$$\mathrm{Rep}(\Gamma) = \{\Phi : \Gamma \to \mathbb{R}^2 \rtimes O(2) \mid \Phi \text{ is a discrete faithful representation}\}$$

2.2.2. Motions in representation space For our purposes, a 1-parameter family of representations is a continuous motion if it is pointwise continuous. More precisely, identify $\mathrm{Euc}(2) \cong \mathbb{R}^2 \times O(2)$ as topological spaces. Suppose $\Phi_t : \Gamma \to \mathrm{Euc}(2)$ is a family of representations defined for $t \in (-\epsilon, \epsilon)$ for some $\epsilon > 0$. Then, Φ_t is a continuous motion through Φ_0 if $\Phi_t(\gamma)$ is a continuous path in $\mathrm{Euc}(2)$ for all $\gamma \in \Gamma$.

2.2.3. Generators for Γ_k To describe the representation space, we need a description of the generating sets for each of the Γ_k, which follows from their descriptions as semi-direct products of $\mathbb{Z}^2 \rtimes (\mathbb{Z}/k\mathbb{Z})$.

Lemma 2.2. *The following are generating sets for each of the Γ_k:*

- Γ_2 *is generated by the set* $\{((1,0),0), ((0,1),0), ((0,0),1)\}$.
- Γ_k *is generated by the set* $\{((1,0),0), ((0,0),1)\}$ *for* $k = 3, 4, 6$.

For convenience, we set the notation $r_k = ((0,0),1)$, $t_1 = ((1,0),0)$, and $t_2 = ((0,1),0)$.

2.2.4. Coordinates for representations We now show how to give convenient coordinates for the representation space for each Γ_k for $k = 2, 3, 4, 6$; by the classification of 2-dimensional crystallographic groups, these are the only cases we need to check. This next lemma follows readily from Bieberbach's Theorems [2, 3] and Lemma 2.2, but we give a proof in Section 2.2.6 for completeness.

Lemma 2.3. *The representation spaces of each of the Γ_k can be given coordinates as follows:*

- $\mathrm{Rep}(\Gamma_2) \cong \{\mathbf{v}_1, \mathbf{v}_2, \mathbf{w} \in \mathbb{R}^2 : \mathbf{v}_1$ *and* $\mathbf{v}_2$ *are linearly independent*$\}$
- $\mathrm{Rep}(\Gamma_k) \cong \{\mathbf{v}_1, \mathbf{w}, \varepsilon \mid \mathbf{v}_1 \neq 0, \varepsilon = \pm 1, \mathbf{v}_1, \mathbf{w} \in \mathbb{R}^2\}$ *for* $k = 3, 4, 6$

The vectors specify the "$\mathbb{R}^2$-part" of the image of a generator in $\mathrm{Euc}(2) \cong \mathbb{R}^2 \rtimes O(2)$. Specifically, $\mathbf{v}_i, \mathbf{w}$ are the coordinates for Φ precisely when $\Phi(t_i) = (\mathbf{v}_i, \mathrm{Id})$ and $\Phi(r_k) = (\mathbf{w}, R_k)$. The vector $\mathbf{w}$ determines the rotation center, but is *not* the rotation center itself. (In fact, the rotation center is $(I - R_k)^{-1}(\mathbf{w})$.)

2.2.5. Coordinates for finite-order rotations The following lemma characterizes an order k rotation in terms of the semidirect product $\mathbb{R}^2 \rtimes O(2)$.

Lemma 2.4. *Let ψ be an orientation-preserving element of* $\mathrm{Euc}(2)$. *Then ψ has order $k = 2, 3, 4, 6$ if and only if it is of the form $(\mathbf{w}, R_k^{\pm 1})$, where R_k is the order k counterclockwise rotation through angle $2\pi/k$ and $\mathbf{w} \in \mathbb{R}^2$.*

Proof. By Lemma 2.1, ψ is a rotation or a translation, and translations clearly have the form $(\mathbf{w}, \mathrm{Id})$. Thus, ψ is a rotation if and only if it has the form $(\mathbf{w}, R)$ for some nontrivial rotation R. If ψ^k has the form $(\mathbf{w}', \mathrm{Id})$, then $\mathbf{w}'$ is necessarily zero as no power of a rotation is a translation. Hence, the order of $(\mathbf{w}, R)$ is precisely that of R and the rest of the theorem follows easily. □

2.2.6. Proof of Lemma 2.3 We let $\Phi \in \mathrm{Rep}(\Gamma_k)$ be a discrete, faithful representation. Thus Φ is determined by the images of the generators, so Lemma 2.2 tells us we need only to check t_1, t_2, and r_k.

The generators t_i must always be mapped to translations: since they are infinite order and Φ is faithful, the only other possibility is an infinite order rotation. This would contradict Φ being discrete. Thus:

- For $k = 2$, the elements t_1 and t_2 are mapped to translations $(\mathbf{v}_1, \mathrm{Id})$ and $(\mathbf{v}_2, \mathrm{Id})$.
- For $k = 3, 4, 6$, the element t_1 is mapped to a translation $(\mathbf{v}_1, \mathrm{Id})$.

Moreover, faithfulness and discreteness force:

- All the images $\mathbf{v}_i$ to be non-zero.
- The images $\mathbf{v}_1$ and $\mathbf{v}_2$ to be linearly independent for $k = 1, 2$.

By Lemma 2.4 we must have $\Phi(r_k) = (\mathbf{w}, R_k^{\varepsilon})$ for some $w \in \mathbb{R}^2$ and $\varepsilon \in \{-1, 1\}$. Since R_2 is order 2, we have $\Phi(r_2) = (\mathbf{w}, R_2) = (\mathbf{w}, R_2^{-1})$, and so ε is unnecessary for Γ_2.

In the other direction, given the data described in the statement of the lemma, we simply define $\Phi(t_i)$ and $\Phi(r_k)$ as above. When $k = 3, 4, 6$, we set $\Phi(t_2) = (R_k^{\varepsilon}\mathbf{v}_1, \mathrm{Id})$. For arbitrary elements of Γ, we define $\Phi((m_1, m_2), m_3) = \Phi(t_1)^{m_1}\Phi(t_2)^{m_2}\Phi(r_k)^{m_3}$. It is straightforward to check Φ as defined is a homomorphism, and that it is discrete and faithful. □

2.2.7. Degenerate representations When we are dealing with "collapsed realizations" of direction networks in Section 4, we will need to work with certain degenerate representations of Γ_k. The space

$$\overline{\mathrm{Rep}}(\Gamma_k)$$

is defined to be representations of Γ_k where we allow the $\mathbf{v}_i$ to be any vectors. Topologically this is the closure of $\mathrm{Rep}(\Gamma_k)$ in the space of all (not necessarily discrete or faithful) representations $\Gamma_k \to \mathrm{Euc}(2)$.

2.2.8. Rotations and translations in crystallographic groups As we have defined them, 2-dimensional crystallographic groups are abstract groups admitting a discrete faithful representation to Euc(2). However, as we saw in the proof of Lemma 2.3, all group elements in $\Lambda(\Gamma_k)$ must be mapped to translations, and all group elements outside $\Lambda(\Gamma_k)$ must be mapped to rotations. Consequently, we will henceforth call elements of $\Lambda(\Gamma_k)$ *"translations"* and elements outside of $\Lambda(\Gamma_k)$ *"rotations"* (even though technically they are elements of Γ_k, not Euc(2)).

2.3 Subgroup structure

This short section contains some useful structural lemmas about subgroups of Γ_k.

2.3.1. The translation subgroup For a subgroup $\Gamma' < \Gamma_k$, we define its *translation subgroup* $\Lambda(\Gamma')$ to be $\Gamma' \cap \Lambda(\Gamma_k)$. (Recall that $\Lambda(\Gamma_k)$ is the subgroup $\mathbb{Z}^2$ coming from the semidirect product decomposition of Γ_k.)

2.3.2. Facts about subgroups With all the definitions in place, we state several lemmas about subgroups of Γ_k that we need later.

Lemma 2.5. *Let $\Gamma' < \Gamma_k$ be a subgroup of Γ_k, and suppose $\Gamma' \neq \Lambda(\Gamma')$. Then Γ' is generated by one rotation and $\Lambda(\Gamma')$.*

Proof. We need only observe that $\Gamma_k/\Lambda(\Gamma_k)$ is finite cyclic and contains $\Gamma'_k/\Lambda(\Gamma'_k)$ as a subgroup. □

This next lemma is straightforward, but useful. We omit the proof.

Lemma 2.6. *Let $r_1, r_2 \in \Gamma_k$ be rotations. Then $\langle r_1, r_2 \rangle$ is a finite cyclic subgroup consisting of rotations if and only if some nontrivial powers r_1^p and r_2^q commute.*

Lemma 2.7. *Let $r' \in \Gamma_2$ be a rotation and $\Gamma' < \Lambda(\Gamma_2)$ a subgroup of the translation subgroup of Γ_2. Then $\Lambda(\langle r', \Gamma' \rangle) = \Gamma'$; i.e., after adding the rotation r', the translation subgroup of the group generated by r' and Γ' is again Γ'.*

Proof. All translation subgroups of Γ_2 are normal, and so the set $\{gh \mid g \in \{r', \mathrm{Id}\} \ \ h \in \Gamma'\}$ is a subgroup and is equal to $\langle r', \Gamma' \rangle$. Clearly, the only translations are those elements of Γ'. □

2.4 The restricted representation space and its dimension

To define our degree of freedom heuristics in Section 3, we need to understand how representations of Γ_k restrict to subgroups $\Gamma' < \Gamma_k$, or equivalently, which representations of Γ' extend to Γ_k. For $\Gamma' < \Gamma_k$, the *restricted representation space* of Γ' is the image of the restriction map from $\mathrm{Rep}(\Gamma_k)$ to $\mathrm{Rep}(\Gamma')$, i.e.,

$$\mathrm{Rep}_{\Gamma_k}(\Gamma') = \{\Phi : \Gamma' \to \mathrm{Euc}(2) \mid \Phi \text{ extends to a discrete faithful representation of } \Gamma_k\}$$

We define the notation $\mathrm{rep}_{\Gamma_k}(\Gamma') := \dim \mathrm{Rep}_{\Gamma_k}(\Gamma')$, since the dimension of $\mathrm{Rep}_{\Gamma_k}(\Gamma')$ is an important quantity in what follows. We also define the following invariant, which is essential for defining our combinatorial matroids:

$$T(\Gamma') := \begin{cases} 0 & \text{if } \Gamma' \text{ has a rotation} \\ 2 & \text{if } \Gamma' \text{ has no rotations} \end{cases}$$

Equivalently, we may define $T(\Gamma')$ as the dimension of the space of translations commuting with $\Phi(\Gamma')$ for any $\Phi \in \mathrm{Rep}(\Gamma_k)$. In Section 4.3, we will show that $T(\Gamma')$ is the dimension of the space of collapsed solutions of a direction network for a connected graph associated with the subgroup Γ'

We now develop some properties of $\mathrm{rep}_{\Gamma_k}(\cdot)$ and how it changes as new generators are added to a finitely generated subgroup. These will be important for counting the degrees of freedom in direction networks (Section 4).

2.4.1. Translation subgroups For translation subgroups $\Gamma' < \Gamma_k$, we are interested in the dimension of $\mathrm{Rep}_{\Gamma_k}(\Gamma')$. The following lemma gives a characterization for translation subgroups in terms of the rank of Γ'.

Lemma 2.8. *Let $\Gamma' < \Gamma_k$ be a* nontrivial *subgroup of translations.*

- *If $k = 3, 4, 6$, then* $\mathrm{rep}_{\Gamma_k}(\Gamma') = 2$.
- *If $k = 2$, then* $\mathrm{rep}_{\Gamma_k}(\Gamma') = 2 \cdot r$*, where r is the rank of* Γ'.

In particular, $\mathrm{rep}_{\Gamma_k}(\Gamma')$ *is even.*

Proof. Suppose $k = 3, 4$ or 6. By Lemma 2.3, the space of representations of Γ_k is 4-dimensional and is uniquely determined by the parameters $\mathbf{v}_1, \mathbf{w}$ and the sign ε. The group $\Lambda(\Gamma_k) \cong \mathbb{Z}^2$ is generated by t_1 and $r_k t_1 r_k^{-1}$, and so any $\gamma \in \Lambda(\Gamma_k)$ can be written uniquely as $t_1^{m_1} r_k t_1^{m_2} r_k^{-1}$ for integers m_1, m_2. Thus, since $\Phi(\gamma)$ is a translation,

$$\begin{aligned}\Phi(\gamma) &= \Phi(t_1)^{m_1}\Phi(r_k)\Phi(t_1)^{m_2}\Phi(r_k)^{-1} \\ &= \left(m_1\mathbf{v}_1, \mathrm{Id}\right)\left(\mathbf{w}, R_k^{\varepsilon}\right)\left(m_2\mathbf{v}_1, \mathrm{Id}\right)\left(\mathbf{w}, R_k^{-\varepsilon}\right) \\ &= \left(m_1\mathbf{v}_1 + m_2 R_k^{\varepsilon}\mathbf{v}_1, \mathrm{Id}\right)\end{aligned}$$

The computation shows that the restriction of Φ to $\Lambda(\Gamma_k)$ is independent of the parameter $\mathbf{w}$. Moreover two representations with the same parameter ϵ restrict to the same representation of $\Lambda(\Gamma_k)$ precisely when the $\mathbf{v}_1$ parameters are equivalent, and $\mathbf{v}_1$ is completely determined by $\Phi(\gamma)$.

Suppose $k = 2$. In this case by the proof of Lemma 2.3, any discrete faithful representation $\Lambda(\Gamma_k) \to \mathrm{Euc}(2)$ extends to a discrete faithful representation of Γ_k. Since $\Lambda(\Gamma_k) \cong \mathbb{Z}^2$, any discrete faithful representation of its subgroups to $\mathbb{R}^2$ extends to $\Lambda(\Gamma_k)$ and hence to Γ_k. Hence $\mathrm{rep}_{\Gamma_k}(\Gamma')$ is equal to the dimension of representations $\Gamma' \to \mathbb{R}^2$ which is twice the rank of Γ'. $\square$

2.4.2. The r-closure of a subgroup In Section 2.6, we will introduce a matroid on the elements of a crystallographic group. To prove the required properties, we need to know how the translation subgroup $\Lambda(\cdot)$ changes as generators are added to a subgroup of Γ_k. We define the *r-closure,* $\mathrm{cl}(\Gamma')$*, of* Γ' to be the largest subgroup containing Γ' such that

$$\mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma')) = \mathrm{rep}_{\Gamma_k}(\Lambda(\mathrm{cl}(\Gamma'))) \quad \text{and} \quad T(\Gamma') = T(\mathrm{cl}(\Gamma')) \tag{5}$$

The letter r in this terminology refers to the rank function r defined in Section 2.6, and the r-closure is defined such that $\mathrm{cl}(\Gamma')$ is the largest subgroup containing Γ' with $r(\mathrm{cl}(\Gamma')) = r(\Gamma')$. The properties of the r-closure are needed to study the matroid defined by the closely related rank function g_1 (also in Section 2.6), which is a building block for the definition of Γ-Laman graphs in Section 3.4. Since there will be no confusion, we will henceforth drop the r and simply refer to closures of subgroups.

2.4.3. Properties of the closure This next sequence of lemmas enumerates the properties of the closure that we will use in the sequel.

Lemma 2.9. *Let $\Gamma' < \Gamma_k$ be a subgroup of Γ_k. Then, the closure* $\mathrm{cl}(\Gamma')$ *is well-defined. Specifically for $k = 2$,*

- *If Γ' is a translation subgroup, then* $\mathrm{cl}(\Gamma')$ *is the subgroup of translations with a non-trivial power in Γ'.*
- *If Γ' has translations and rotations, then* $\mathrm{cl}(\Gamma') = \langle r', \mathrm{cl}(\Lambda(\Gamma'))\rangle$ *for any rotation $r' \in \Gamma'$.*

For $k = 3, 4, 6$, there are four possibilities for the closure:

- *If Γ' is trivial, then the closure is trivial.*
- *If Γ' is cyclic, then the closure is a cyclic subgroup of order k.*
- *If Γ' is a nontrivial translation subgroup, then the closure is the translation subgroup of Γ_k.*
- *If Γ' has translations and rotations, then the closure is all of Γ_k.*

Proof. First let $k = 2$. There are two cases. If Γ' contains only translations, we set

$$\mathrm{cl}(\Gamma') = \{t \in \Lambda(\Gamma_2) : t^i \in \Gamma' \text{ for some power } i \text{ of } t\}$$

Any subgroup $\Gamma'' < \Gamma_2$ containing Γ' with $T(\Gamma') = T(\Gamma'')$ and $\mathrm{rep}_{\Gamma_k}(\Gamma') = \mathrm{rep}_{\Gamma_k}(\Gamma'')$ must be a translation group of the same rank as Γ' and $\mathrm{cl}(\Gamma')$ is the largest such subgroup.

Otherwise, Γ' contains a rotation r'. In this case, we set

$$\mathrm{cl}(\Gamma') = \langle r', \mathrm{cl}(\Lambda(\Gamma'))\rangle$$

By Lemma 2.7, for $\mathrm{cl}(\Gamma')$ defined this way, the translation subgroup $\Lambda(\mathrm{cl}(\Gamma'))$ is just $\mathrm{cl}(\Lambda(\Gamma'))$ which by the previous paragraph is the largest translation subgroup containing $\Lambda(\Gamma')$ and having the same rank. Suppose $\Gamma'' < \Gamma_2$ contains Γ' and satisfies $\mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma'')) = \mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma'))$ and $T(\Gamma'') = T(\Gamma')$. Then, $\Gamma'' = \langle r', \Lambda(\Gamma'')\rangle$ and $\Lambda(\Gamma')$ and $\Lambda(\Gamma'')$ have the same rank. This implies that $\Lambda(\Gamma'') < \mathrm{cl}(\Lambda(\Gamma'))$ and thus $\Gamma'' < \mathrm{cl}(\Gamma')$.

Now we suppose that $k = 3, 4, 6$. There are four possibilities for Γ':

- If Γ' is trivial, then we define $\mathrm{cl}(\Gamma') = \Gamma'$.
- If Γ' is a cyclic group of rotations, then Lemma 2.6 guarantees that there is a unique largest cyclic subgroup containing it, and we define this to be $\mathrm{cl}(\Gamma')$. Any larger group will have a different rep_{Γ_k} value.
- If Γ' has only translations, then we define $\mathrm{cl}(\Gamma') = \Lambda(\Gamma_k)$. From Lemma 2.8 it follows that $\mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma')) = \mathrm{rep}_{\Gamma_k}(\Lambda(\mathrm{cl}(\Gamma')))$. Any larger subgroup will have a different $T(\cdot)$ value.
- If Γ' has translations and rotations, then it has the same $\mathrm{rep}_{\Gamma_k}(\Lambda(\cdot))$ and $T(\cdot)$ values as Γ_k, so $\mathrm{cl}(\Gamma') = \Gamma_k$.

□

Lemma 2.10. *Let $\Gamma' < \Gamma_k$ be a finitely-generated subgroup of Γ_k, and let $\Gamma'' < \Gamma'$ be a subgroup of Γ'. Then $\mathrm{cl}(\Gamma'') < \mathrm{cl}(\Gamma')$.*

Proof. Pick a generating set of Γ'' that extends to a generating set of Γ'. Analyzing the cases in Lemma 2.9 shows that the closure cannot become smaller after adding generators. □

Lemma 2.11. *Let $\Gamma' < \Gamma_k$ be a translation subgroup of Γ_k, and let $\gamma \in \Gamma_k$. Then $\mathrm{cl}(\gamma\Gamma'\gamma^{-1}) = \mathrm{cl}(\Gamma')$; i.e., the closure of translation subgroups is fixed under conjugation.*

Proof. For $k = 2$ this follows from the fact that all translation subgroups are normal. For $k = 3, 4, 6$ it is immediate from Lemma 2.9. □

Lemma 2.12. *Let $\Gamma' < \Gamma_k$ be a subgroup of Γ_k, and let $\Gamma'' < \Lambda(\Gamma_k)$ be a translation subgroup of Γ_k. Then $\mathrm{cl}(\langle\Lambda(\Gamma'), \Gamma''\rangle) = \mathrm{cl}(\Lambda(\langle\Gamma', \Gamma''\rangle))$.*

Proof. The proof is in cases based on k. For $k = 3, 4, 6$, Lemma 2.9 implies that either Γ'' is trivial or both sides of the desired equation are $\Lambda(\Gamma_k)$. In either case, the lemma follows at once.

Now suppose that $k = 2$. If Γ' is a translation subgroup, then the lemma follows immediately. Otherwise, we know that Γ' is generated by a rotation r' and the translation subgroup $\Lambda(\Gamma')$. Applying Lemma 2.7, we see that

$$\Lambda(\langle\Gamma', \Gamma''\rangle) = \Lambda(\langle r', \Lambda(\Gamma'), \Gamma''\rangle) = \langle\Lambda(\Gamma'), \Gamma''\rangle$$

from which the lemma follows. □

2.4.4. The quantity $\mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma')) - T(\Gamma')$ The following statement plays a key role in the matroidal construction of Section 2.6.

Proposition 2.13. *Let $\Gamma' < \Gamma_k$ be a subgroup of Γ_k, and let $\gamma \in \Gamma_k$ be an element of Γ_k. Then,*

$$\mathrm{rep}_{\Gamma_k}(\Lambda(\langle\Gamma', \gamma\rangle)) - T(\langle\Gamma', \gamma\rangle) - (\mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma')) - T(\Gamma')) = \begin{cases} 2 & \textit{if } \gamma \notin \mathrm{cl}(\Gamma') \\ 0 & \textit{otherwise} \end{cases}$$

i.e., the quantity $\mathrm{rep}_{\Gamma_k}(\Lambda(\cdot)) - T(\cdot)$ increases by two after adding γ to Γ' if and only if $\gamma \notin \mathrm{cl}(\Gamma')$ and otherwise the increase is zero.

Proof. If $\gamma \in \mathrm{cl}(\Gamma')$, this follows at once from the definition, since the quantity $\mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma')) - T(\Gamma')$ depends only on the closure.

Now suppose that $\gamma \notin \mathrm{cl}(\Gamma')$. Since the closure is defined in terms of $\mathrm{rep}_{\Gamma_k}(\Lambda(\cdot))$ and $T(\cdot)$, Lemma 2.10 implies that at least one of $\mathrm{rep}_{\Gamma_k}(\Lambda(\cdot))$ or $-T(\cdot)$ increases. It is easy to see from the definition that either type of increase is by at least 2. We will show that the increase is at most 2, from which the lemma follows. The rest of the proof is in three cases, depending on k.

Now we let $k = 3, 4, 6$. The only way for the increase to be larger than 2 is for Γ' to be trivial and $\mathrm{cl}(\langle\gamma\rangle) = \Gamma_k$. This is impossible given the description from Lemma 2.9.

To finish, we address the case $k = 2$. Suppose γ is a translation. Then $T(\langle\gamma, \Gamma'\rangle) = T(\langle\Gamma'\rangle)$, since adding γ as a generator does not give us a new rotation if one was not already present in Γ'. Lemmas 2.5 and 2.7 imply that $\Lambda(\langle\gamma, \Gamma'\rangle) = \langle\gamma, \Lambda(\Gamma')\rangle$. Hence, the rank of the translation subgroup increases by at most 1, and so, by Lemma 2.8, $\mathrm{rep}_{\Gamma_k}(\cdot)$ increases by at most 2.

Now suppose that γ is a rotation. If Γ' has no rotations, then Lemma 2.7 implies $\Lambda(\langle\gamma, \Gamma'\rangle) = \Gamma'$, and so $T(\cdot)$ decreases and $\mathrm{rep}_{\Gamma_2}(\cdot)$ is unchanged. If Γ' has rotations, then $\Gamma' = \langle r', \Lambda(\Gamma')\rangle$ for some rotation $r' \in \Gamma'$. Since $k = 2$, the product $r'\gamma$ is a translation and so

$$\Lambda(\langle\gamma, \Gamma'\rangle) = \Lambda(\langle\gamma, r', \Lambda(\Gamma')\rangle) = \Lambda(\langle r', r'\gamma, \Lambda(\Gamma')\rangle) = \langle r'\gamma, \Lambda(\Gamma')\rangle$$

Thus, in this case, the the number of generators of the translation subgroup increases by at most one and $T(\cdot)$ is unchanged. By Lemma 2.8, the proof is complete. □

2.5 Teichmüller space and the centralizer

The representation spaces defined in the previous two sections are closely related to the degrees of freedom in the crystallographic direction networks we study in the sequel. In this section, we discuss the *Teichmüller space* and *centralizer*, which play the same role for frameworks.

2.5.1. Teichmüller space The *Teichmüller space* of Γ_k is defined to be the space of discrete faithful representations, modulo conjugation by Euc(2); i.e. $\mathrm{Teich}(\Gamma_k) = \mathrm{Rep}(\Gamma_k)/\mathrm{Euc}(2)$. For a subgroup $\Gamma' < \Gamma_k$, we define its *restricted Teichmüller space* to be

$$\mathrm{Teich}_{\Gamma_k}(\Gamma') = \mathrm{Rep}_{\Gamma_k}(\Gamma')/\mathrm{Euc}(2)$$

Correspondingly, we define $\mathrm{teich}_{\Gamma_k}(\Gamma') = \dim(\mathrm{Teich}_{\Gamma_k}(\Gamma'))$.

2.5.2. The centralizer For a subgroup $\Gamma' \le \Gamma_k$ and a discrete faithful representation $\Phi : \Gamma_k \to \mathrm{Euc}(2)$, the *centralizer of* $\Phi(\Gamma')$ which we denote $\mathrm{Cent}_{\mathrm{Euc}(2)}(\Phi(\Gamma'))$ is the set of elements commuting with all elements in $\Phi(\Gamma')$. We define $\mathrm{cent}(\Gamma')$ to be the dimension of the centralizer $\mathrm{Cent}_{\mathrm{Euc}(2)}(\Phi(\Gamma'))$. The quantity $\mathrm{cent}(\Gamma')$ is independent of Φ, and we can compute it. Since we do not depend on Lemma 2.14 or Proposition 2.15 for any of our main results, we skip the proofs in the interest of space.

Lemma 2.14. *Let notation be as above. The quantity* $\mathrm{cent}(\Gamma')$ *is independent of the representation* Φ. *Furthermore,* $\mathrm{cent}(\Gamma') \ge T(\Gamma')$*, and in particular,*

$$\mathrm{cent}(\Gamma') = \begin{cases} 0 & \text{if } \Gamma' \text{ contains rotations and translations} \\ 1 & \text{if } \Gamma' \text{ contains only rotations} \\ 2 & \text{if } \Gamma' \text{ contains only translations} \\ 3 & \text{if } \Gamma' \text{ is trivial} \end{cases}$$

As a corollary, we get the following proposition relating $\mathrm{rep}_{\Gamma_k}(\cdot)$ and $T(\cdot)$ to $\mathrm{teich}_{\Gamma_k}(\cdot)$ and $\mathrm{cent}(\cdot)$.

Proposition 2.15. *Let $\Gamma' < \Gamma_k$. Then:*

- ***(A)*** *If Γ' contains a translation, then $T(\Gamma') = \mathrm{cent}(\Gamma')$. Otherwise, $T(\Gamma') = \mathrm{cent}(\Gamma') - 1$.*
- ***(B)*** *If Γ' is a non-trivial translation subgroup, then* $\mathrm{teich}_{\Gamma_k}(\Gamma') = \mathrm{rep}_{\Gamma_k}(\Gamma') - 1$.
- ***(C)*** *If Γ' is trivial, then* $\mathrm{teich}_{\Gamma}(\Gamma') = \mathrm{rep}_{\Gamma}(\Gamma') = 0$.
- ***(D)*** *For any $\Gamma' < \Gamma_k$,* $\mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma')) - T(\Gamma') = \mathrm{teich}_{\Gamma_k}(\Lambda(\Gamma')) - \mathrm{cent}(\Gamma')$.

2.6 A matroid on crystallographic groups

We now define and study a matroid $M_{\Gamma_k,n}$ for $k = 2,3,4,6$.

2.6.1. Preview of Γ-$(1,1)$ graphs and $M_{\Gamma_k,n}$ In Section 3.5, we will relate $M_{\Gamma_k,n}$ to "Γ-$(1,1)$ graphs", which are defined in Section 3.3.6. The results here, roughly speaking, are the group theoretic part of the proof of Proposition 3.5 in Section 3.5.

We briefly motivative the definitions given next. In general, Γ-$(1,1)$ graphs need not be connected, and each connected component has an associated finitely generated subgroup of Γ_k. The ground set of $M_{\Gamma_k,n}$ and the A_i defined below capture this situation. The operations of conjugating and fusing, defined here in Sections 2.6.9 and 2.6.10 will be interpreted graph theoretically in Section 3.5.

2.6.2. The ground set For the definition of the ground set, we fix Γ_k and a natural number $n \geq 1$. The ground set $E_{\Gamma_k,n}$ is defined to be:

$$E_{\Gamma_k,n} = \{(\gamma, i) : 1 \leq i \leq n\}$$

In other words the ground set is n labeled copies of Γ_k.

Let $A \subset E_{\Gamma_k,n}$. We define some notation:

- $A_i = \{\gamma : (\gamma, i) \in A\}$; i.e., A_i is the group elements from copy i of Γ_k in A. Some of the A_i may be empty and A_i can be a multi-set. A may equivalently be defined by the A_i.
- $\Gamma_{A,i} = \langle \gamma : \gamma \in A_i \rangle$; i.e., the subgroup generated by the elements in A_i.
- $\Lambda(A) = \langle \Lambda(\Gamma_{A,1}), \Lambda(\Gamma_{A,2}), \ldots, \Lambda(\Gamma_{A,n}), \rangle$; the translation subgroup generated by the translations in each of the $\Gamma_{A,i}$.
- $c(A)$ is the number of A_i that are not empty.

2.6.3. The rank function We now define the function $g_1(A)$ for $A \subset E_{\Gamma_k,n}$ to be

$$g_1(A) = n + \frac{1}{2}\operatorname{rep}_{\Gamma_k}(\Lambda(A)) - \frac{1}{2}\sum_{i=1}^{n} T(\Gamma_{A,i})$$

The meaning of the terms in $g_1(A)$ are as follows:

- The second term is a global adjustment for the representation space of the group generated by the translations in each of the $\Gamma_{A,i}$. We note that this is *not* the same as the translation group $\Lambda(\langle \gamma : \gamma \in \cup_{i=1}^{n} A_i)$, which includes translations arising as products of rotations in different A_i.
- The quantity $n - \frac{1}{2}\sum_{i=1}^{n} T(\Gamma_{A,i}) = \sum_{i=1}^{n}(1 - \frac{1}{2}T(\Gamma_{A,i}))$ is a local adjustment based on whether $\Gamma_{A,i}$ contains a rotation: each term in the latter sum is one if $\Gamma_{A,i}$ contains a rotation and otherwise it contributes nothing.

2.6.4. An analogy to uniform linear matroids To give some intuition about why the construction above might be matroidal, we observe that Proposition 2.13, interpreted in matroidal language gives us:

Proposition 2.16. *Let A be a finite subset of Γ_k generating a subgroup Γ_A. Then the function*

$$r(A) = \frac{1}{2}\left(\operatorname{rep}_{\Gamma_k}(\Lambda(\Gamma_A)) - T(\Gamma_A)\right)$$

is the rank function of a matroid on the ground set Γ_k.

The matroid in the conclusion of Proposition 2.16 is analogous to a linear matroid, with Γ_k playing the role of a vector space and r the role of dimension of the linear span. (And, in fact, for the group $\mathbb{Z}^2$, r reduces simply to linear independence, as in [16, Section 4].) Since the function g_1, defined above, builds on r, one might expect that it inherits a matroidal structure. We verify this next.

2.6.5. $M_{\Gamma_k,n}$ is a matroid The following proposition is the main result of this section.

Proposition 2.17. *The function g_1 is the rank function of a matroid $M_{\Gamma_k,n}$.*

We note that although the ground set is infinite, since our matroids are finite rank, all the facts for finite matroids which we cite apply here as well.

The proof depends on Lemmas 2.18 and 2.19 below, so we defer it for the moment to Section 2.6.6. The strategy is based on the observation that when $n = 1$, the ground set is essentially Γ_k. In this case, submodularity and normalization of g_1 (the most difficult properties to establish) follow immediately from Proposition 2.13. The motivation of Lemmas 2.18 and 2.19 is to reduce, as much as possible, the proof of the general case to $n = 1$.

Lemma 2.18. *Let $A \subset E_{\Gamma_k,n}$, and set $\Gamma'_{A,\ell} = \langle \Gamma_{A,\ell}, \Lambda(A)\rangle$. Then, for all $1 \le \ell \le n$,*

- $\operatorname{cl}(\Lambda(A)) = \operatorname{cl}(\Lambda(\Gamma'_{A,\ell}))$
- $T(\Gamma_{A,\ell}) = T(\Gamma'_{A,\ell})$

Proof. The first statement is immediate from Lemma 2.12. The second statement follows from the fact that $\Lambda(A)$ is a translation subgroup of Γ_k, so $\Gamma'_{A,\ell}$ has a rotation if and only if $\Gamma_{A,\ell}$ does. $\square$

Lemma 2.19. *Let $A \subset E_{\Gamma_k,n}$, and set $\Gamma'_{A,\ell} = \langle \Gamma_{A,\ell}, \Lambda(A)\rangle$. If $B = A + (\gamma, \ell)$ and $\Gamma'_{B,\ell} = \langle \Gamma_{B,\ell}, \Lambda(B)\rangle$, Then,*

$$\Gamma'_{B,\ell} = \langle \gamma, \Gamma'_{A,\ell}\rangle$$

Proof. First we observe that

$$\Gamma'_{B,\ell} = \langle \gamma, \Gamma_{A,\ell}, \Lambda(B)\rangle \geq \langle \gamma, \Gamma_{A,\ell}, \Lambda(A)\rangle$$

so to finish the proof we just have to show that

$$\Lambda(B) \leq \langle \gamma, \Gamma_{A,\ell}, \Lambda(A)\rangle$$

Since $\Gamma_{B,i} = \Gamma_{A,i}$ for all $i \neq \ell$, it follows that

$$\Lambda(B) = \langle \Lambda(\langle \gamma, \Gamma_{A,\ell}\rangle)\,,\ \Lambda(A)\rangle \leq \langle \gamma, \Gamma_{A,\ell}, \Lambda(A)\rangle$$

$\square$

2.6.6. Proof of Proposition 2.17 We check the rank function axioms [17, Section 1.3].

Non-negativity: This follows from the fact that $\mathrm{rep}_{\Gamma_k}(\cdot)$ is non-negative, and the sum of the $\frac{1}{2}T(\cdot)$ terms cannot exceed n.

Monotonicity: This is immediate from Lemma 2.10 and the fact that $-T(\cdot)$ and $\mathrm{rep}_{\Gamma_k}(\Lambda(\cdot))$ only increase when the size of the closure increases.

Normalization: To prove that g_1 is normalized, let $A \subset E_{\Gamma_k,n}$ and $B = A + (\gamma, \ell)$. Since all the $T(\Gamma'_{\cdot,i})$ terms cancel except for the ones with $i = \ell$, the increase is given by

$$g_1(B) - g_1(A) = \frac{1}{2}\left(\mathrm{rep}_{\Gamma_k}(\Lambda(B)) - \mathrm{rep}_{\Gamma_k}(\Lambda(A)) - T(\Gamma_{B,\ell}) + T(\Gamma_{A,\ell})\right)$$

Because the r.h.s. is an invariant of the closure by Proposition 2.13, we pass to closures and apply Lemma 2.18 to see that the r.h.s. is equal to

$$\frac{1}{2}\left(\mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma'_{B,\ell})) - \mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma'_{A,\ell})) - T(\Gamma'_{B,\ell}) + T(\Gamma'_{A,\ell})\right)$$

Using Lemma 2.19 then tells us that this can be simplified further to

$$\frac{1}{2}\left(\mathrm{rep}_{\Gamma_k}(\Lambda(\langle \gamma, \Gamma'_{A,\ell}\rangle)) - \mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma'_{A,\ell})) - T(\langle \gamma, \Gamma'_{A,\ell}\rangle) + T(\Gamma'_{A,\ell})\right)$$

at which point Proposition 2.13 applies, and we conclude that the increase is either zero or one.

Submodularity: We will verify the following form of the submodular inequality:

$$f(A \cup \{(\gamma, \ell)\}) - f(A) \geq f(B \cup \{(\gamma, \ell)\}) - f(B) \qquad \text{for all } A \subset B \tag{6}$$

Inspecting the argument for normalization and Proposition 2.13, we see that the r.h.s., is positive only if $\gamma \notin \mathrm{cl}(\Gamma_{B,\ell})$), in which case it is always 2. By Lemma 2.10, for this γ, we also have $\gamma \notin \mathrm{cl}(\Gamma_{A,\ell})$, so the l.h.s. is also 2. Because both sides are always non-negative, (6) follows. $\square$

2.6.7. The bases and independent sets With the rank function of $M_{\Gamma_k,n}$ determined, we can give a structural characterization of its bases and independent sets. Let $A \subset E_{\Gamma_k,n}$. We define A to be *independent* if

$$|A| = g_1(A)$$

If A is independent and, in addition

$$|A| = c(A) + \frac{1}{2}\,\mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma_k))$$

we define A to be *tight*. A (not-necessarily independent) set A with $c(A)$ parts that contains a tight subset on $c(A)$ is defined to be *spanning*.

We define the classes

$$\begin{aligned}
\mathcal{B}(M_{\Gamma_k,n}) &= \left\{B \subset E_{\Gamma_k,n} : B \text{ is independent and } |B| = n + \frac{1}{2}\,\mathrm{rep}_{\Gamma_k}(\Gamma_k)\right\} \\
\mathcal{I}(M_{\Gamma_k,n}) &= \left\{B \subset E_{\Gamma_k,n} : B \text{ is independent}\right\}
\end{aligned}$$

It is now immediate from Proposition 2.17 that:

Lemma 2.20. *The classes $\mathcal{I}(M_{\Gamma_k,n})$ and $\mathcal{B}(M_{\Gamma_k,n})$ are the independent sets and bases of the matroid $M_{\Gamma_k,n}$.*

2.6.8. Structure of tight sets We also have a structural characterization of the tight independent sets in $M_{\Gamma_k,n}$.

Lemma 2.21. *An independent set $A \in \mathcal{I}(M_{\Gamma_k,n})$ is tight if and only if it is one of two types:*

- ***(A)*** *Each of the non-empty A_i contains a rotation. One exceptional non-empty A_i contains $\frac{1}{2}\,\mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma_k))$ additional elements, and $\mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma_{A,i})) = \mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma_k))$, and all the rest of the A_i contain a single rotation only.*

- ***(B)*** *Each of the non-empty A_i contains a rotation. Two exceptional non-empty A_i (w.l.o.g., A_1 and A_2) contain, between them, $\frac{1}{2}\,\mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma_k))$ additional elements and*

$$\mathrm{rep}_{\Gamma_k}(\langle \Lambda(\Gamma_{A,1}), \Lambda(\Gamma_{A,2})\rangle) = \mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma_k)).$$

*Type **(B)** is only possible when $\Gamma_k = \Gamma_2$.*

Proof. One direction is straightforward: A set $A \subset E_{\Gamma_k,n}$ of either type **(A)** or **(B)** satisfies, by hypothesis, $|A| = c(A) + \frac{1}{2}\,\mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma_k))$; by construction $T(\Gamma_{A,i})$ is zero for all the non-empty A_i and $\mathrm{rep}_{\Gamma_k}(\Lambda(A)) = \mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma_k))$.

On the other hand, assuming that A is tight, we see that each non-empty part has to contain a rotation, and, since A is independent there are only one (for $k = 3,4,6$) or two ($k = 2$) additional elements in A. For $k = 3,4,6$, the single A_i containing the extra element must generate a translation in which case $\mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma_{A,i})) = \mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma_k))$. For $k = 2$, the translation subgroups of the A_i containing the extra elements must generate a rank 2 translation subgroup of $\Lambda(\Gamma_k)$, and the desired conclusion follows. $\square$

2.6.9. Conjugation of independent sets Let $A \in \mathcal{I}(M_{\Gamma_k,n})$ be an independent set, and suppose, w.l.o.g., that $A_1, A_2, \ldots, A_{c(A)}$ are the non-empty parts of A. Let $\gamma_1, \gamma_2, \ldots, \gamma_{c(A)}$ be elements of Γ_k. The *conjugation of A by* $\gamma_1, \gamma_2, \ldots, \gamma_{c(A)}$ is defined to be

$$\left\{(\gamma_i^{-1} A_i \gamma_i, i) : 1 \le i \le c(A)\right\}$$

Lemma 2.22. *Let $A \in \mathcal{I}(M_{\Gamma_k})$ be an independent set. Then the conjugation of A by $c(A)$ elements $\gamma_1, \ldots, \gamma_{c(A)}$ is also independent.*

Proof. Lemma 2.11 implies that the closure of translation subgroups is preserved under conjugation, and whether or not A_i contains a rotation is preserved as well. Since the rank function g_1 is determined by these two properties of the A_i, we are done. □

2.6.10. Separating and fusing independent sets Let $A \in \mathcal{I}(M_{\Gamma_k})$ be an independent set. A *separation of A* is defined to be the following operation:

- Select i and j such that A_j is empty.
- Select a (potentially empty) subset $A_i' \subset A_i$ of A_i.
- Replace all elements $(\gamma, i) \in A_i'$ with (γ, j).

Lemma 2.23. *Let $A \in \mathcal{I}(M_{\Gamma_k})$ be an independent set. Then any separation of A is also an independent set.*

Proof. Let B be a separation of A. If the subset A_i' in the definition of a separation is empty, then B is the same as A, and there is nothing to prove.

An independent set is either tight or a subset of a tight set. (Bases in particular are tight.) Consequently, by Lemma 2.21, either B_i or B_j consists of a single element. Assume w.l.o.g., it is B_j. Define $C \subset E_{\Gamma_k,n}$ as $C_k = B_k$ for $k \neq j$ and C_j empty; i.e. C is B with the single element in B_j dropped. Then C is a subset of A and hence independent. If B_j consists of a rotation, then adding it to C clearly preserves independence. If B_j consists of a translation γ, then, since A is independent, we must have $\gamma \notin \text{cl}(\Lambda(C))$. Consequently $B = C + (\gamma, j)$ is independent since $\text{cl}(\Lambda(B)) \gneqq \text{cl}(\Lambda(C))$ and hence $\text{rep}_{\Gamma_k}(\Lambda(B)) > \text{rep}_{\Gamma_k}(\Lambda(C))$. □

The reverse of separation is *fusing a set A on A_i and A_j*. This operation replaces A_i with $A_i \cup A_j$ and makes A_j empty. Fusing does not, in general, preserve independence, but it takes tight sets to spanning ones.

Lemma 2.24. *Let A be a tight independent set, and suppose that A_i and A_j are non-empty. Then, after fusing A on A_i and A_j, the result is a spanning set (with one less part).*

Proof. Let B be the set resulting from fusing A on A_i and A_j. By hypothesis, all the non-empty A_ℓ contain a rotation, so this is true of the non-empty B_ℓ as well. The lemma then follows by noting that $\Lambda(A) \le \Lambda(B)$, so the same is true of the closures by Lemma 2.10. Thus, $g_1(B) = c(B) + \text{rep}_{\Gamma_k}(\Lambda(\Gamma_k))$, and this implies B is spanning. □

3. Matroidal sparse graphs

3.1 Colored graphs and the map ρ

We will use *colored graphs*, which are also known as "gain graphs" (e.g., [19]) or "voltage graphs" (e.g. [28]) as the combinatorial model for crystallographic frameworks and direction networks. In this section we give the definitions and explain the relationship between colored graphs and graphs with a free Γ_k-action.

3.1.1. Colored graphs Let $G = (V, E)$ be a finite, directed graph, with n vertices and m edges. We allow multiple edges and self-loops, which are treated the same as other edges. A Γ_k*-colored-graph* (shortly, *colored graph*) (G, γ) is a finite, directed multigraph G and an assignment $\gamma = (\gamma_{ij})_{ij\in E(G)}$ of a group element $\gamma_{ij} \in \Gamma_k$ (the "color") to each edge $ij \in E(G)$.

3.1.2. The covering map Although we work with colored graphs because they are technically easier, crystallographic frameworks were defined in terms of infinite graphs $\tilde{G}$ with Γ_k acting freely and with finite quotient by the representation $\varphi : \Gamma_k \to \mathrm{Aut}(\tilde{G})$. In fact, the formalisms are equivalent, via a specialization of covering space theory (e.g., [9, Section 1.3]). We provide the dictionary here for completeness.

Let (G, γ) be a colored graph, we define its *lift* $\tilde{G} = (\tilde{V}, \tilde{E})$ by the following construction:

- For each vertex $i \in V(G)$, there is a subset of vertices $\{i_\gamma\}_{\gamma\in\Gamma} \subset V(\tilde{G})$ (the fiber over i).
- For each (directed) edge $ij \in E(G)$ with color γ_{ij}, and for each $\gamma \in \Gamma_k$, there is an edge $i_\gamma j_{\gamma\cdot\gamma_{ij}}$ in $E(\tilde{G})$ (the fiber over ij).
- The Γ-action on vertices is $\gamma \cdot i_{\gamma'} = i_{\gamma\gamma'}$. The action on edges is that induced by the vertex action.

Now let $(\tilde{G}, \varphi)$ be an infinite graph with a free Γ_k-action that has finite quotient. We associate a colored graph (G, γ) to $(\tilde{G}, \varphi)$ by the following construction, which we define to be a *colored quotient*:

- Let $G = \tilde{G}/\Gamma$ be the quotient of $\tilde{G}$ by Γ, and fix an (arbitrary) orientation of the edges of G to make it a directed graph. By hypothesis, the vertices of G correspond to the vertex orbits in $\tilde{G}$ and the edges to the edge orbits in $\tilde{G}$
- For each vertex orbit under Γ in $\tilde{G}$, select a representative $\tilde{i}$.
- For each edge orbit in $\tilde{G}$ there is a unique edge $\widetilde{ij}$ that has the representative $\tilde{i}$ as its tail. There is also a unique element $\gamma_{ij} \in \Gamma$ such that the head of $\widetilde{ij}$ is $\gamma_{ij}(\tilde{j})$. We define this γ_{ij} to be the color on the edge $ij \in G$.

From the definition, we see that the specific colored quotient depends on the choice of representatives, however they are all related as follows. For any choice of representatives, the lift $\tilde{G}'$ is isomorphic to $\tilde{G}$ as a graph, and this isomorphism is φ-equivariant. It the follows that the lifts of any two colored quotients are isomorphic to each other via a φ-equivariant map.

The *projection map* from $(\tilde{G}, \varphi)$ to its colored quotient is the function that sends a vertex $\tilde{i} \in V(\tilde{G})$ its representative $i \in V(G)$. Figures 1 and 2 both show examples; the color coding of

the vertices in the infinite developments indicate the fibers over vertices in the colored quotient. The discussion above shows:

Lemma 3.1. *Let (G,γ) be a Γ_k-colored graph. Then its lift is well defined, and is an infinite graph with a free Γ_k-action. If $(\tilde{G},\varphi)$ is an infinite graph with a free Γ_k-action, then it is the lift of any of its colored quotient, and the projection map is well-defined and a covering map.*

3.1.3. The map ρ Let (G,γ) be a colored graph, and let $P = \{e_1, e_2, \ldots, e_t\}$ be any *closed path* in G; i.e., P is a (not necessarily simple) walk in G that starts and ends at the same vertex crossing the edges e_i in order. If we select a vertex b as a *base point*, then the closed paths represent elements of the *fundamental group* $\pi_1(G,b)$.

We define the map ρ as:

$$\rho(P) = \gamma_{e_1}^{\epsilon_1} \cdots \gamma_{e_t}^{\epsilon_t}$$

where ϵ_i is 1 if P crosses e_i in the forward direction (from tail to head) and -1 otherwise. For a connected graph G and choice of base vertex i, the map ρ induces a well-defined homomorphism $\rho : \pi_1(G,i) \to \Gamma$.

3.2 The subgroup of a Γ_k-colored graph

The map ρ, defined in the previous section, is fundamental to the results of this paper. In this section, we develop properties of the ρ-image of a colored graph (G,γ) and connect it with the matroid $M_{\Gamma_k,n}$ which was defined in Section 2.6.

3.2.1. Colored graphs with base vertices Let (G,γ) be a colored graph with n vertices and c connected components $G_1, G_2, \ldots, G_c$. We select a *base vertex* b_i in each connected component G_i, and denote the set of base vertices by B. The triple (G,γ,B) is then defined to be a *marked colored graph*.

If (G,γ,B) is a marked colored graph then ρ induces a homomorphism from $\pi_1(G_i,b_i)$ to Γ_k. In the rest of this section, we show how to use these homomorphisms to define a map from (G,γ) to $E_{\Gamma_k,n}$, the ground set of the matroid $M_{\Gamma_k,n}$.

3.2.2. Fundamental closed paths generate the ρ-image Let (G,γ,B) be a marked colored graph with n vertices and c connected components. Select and fix a maximal forest F of G, with connected components $T_1, T_2, \ldots, T_c$. The T_i are spanning trees of the connected components G_i of G, with the convention that when a connected component G_i has no edges there is a one-vertex "empty tree" T_i.

With this data, we define, for each edge $ij \in E(G) - E(F)$ the *fundamental closed path of ij* to be the path that:

- Starts at the base vertex b_ℓ in the same connected component G_ℓ as i and j.
- Travels the unique path in T_ℓ to i.
- Crosses ij.
- Travels the unique path in T_ℓ back to b_ℓ.

Fundamental closed paths with respect to F in G_i generate $\pi_1(G_i,b_i)$ by [9, Proposition 1A.2].

3.2.3. From colored graphs to sets in $E_{\Gamma_k,n}$ We now let (G,γ,B) be a marked colored graph and fix a choice of spanning forest F. We associate with (G,γ,B,F) a subset $A(G,B,F)$ of $E_{\Gamma_k,n}$ (defined in Section 2.6) as follows:

- For each edge $ij \in E(G_\ell) - E(T_\ell)$, let P_{ij} be the fundamental closed path of ij with respect to T_ℓ and b_ℓ.
- Add an element $(\rho(P_{ij}), \ell)$ to $A(G,B,F)$.

The following is immediate from the previous discussion.

Lemma 3.2. *Using the notation from Section 2.6, $\Gamma_{A,\ell} = \rho(\pi_1(G_\ell, b_\ell))$ where $A = A(G,B,F)$.*

Since we will show, in Section 3.3, that the invariants we need are independent of B and F, we frequently suppress them from the notation when the context is clear.

3.3 Γ-$(2,2)$ graphs

In this section we define Γ-$(2,2)$ *graphs* which are the first of two key families of colored graphs introduced in this paper (the second is Γ-*colored-Laman graphs*, defined in Section 3.4). We also state the main combinatorial results on Γ-$(2,2)$ graphs, but defer the proof of a key technical result, Proposition 3.5, to Section 3.5.

3.3.1. The translation subgroup of a colored graph Let (G,γ,B) be a marked colored graph, as in Section 3.2, with connected components $G_1, G_2, \ldots, G_c$ and base vertices $b_1, b_2, \ldots, b_c$. Recall from Section 3.1.3 that, with this data, there are homomorphisms

$$\rho : \pi_1(G_i, b_i) \to \Gamma_k$$

We define $\Lambda(G,B)$ to be

$$\Lambda(G,B) = \langle \Lambda(\rho(G_i, b_i)) : i = 1,2,\ldots,c \rangle$$

We define $\mathrm{rep}_{\Gamma_k}(G) = \mathrm{rep}_{\Gamma_k}(\Lambda(G,B))$. As the notation suggests, $\mathrm{rep}_{\Gamma_k}(G)$ is independent of the choice of base vertices B.

Lemma 3.3. *Let (G,γ,B) be a marked colored connected graph. The quantity $\mathrm{rep}_{\Gamma_k}(G)$ is independent of the choice of base vertices, and so is a property of the underlying colored graph (G,γ).*

Proof. Changing base vertices corresponds to conjugation. Lemma 2.11 implies that the closure of $\Lambda(G,B)$ is preserved under conjugation. Since $\mathrm{rep}_{\Gamma_k}(\cdot)$ depends only on the closure, the lemma follows. $\square$

3.3.2. The quantity T for a colored graph Let (G,γ,B) be a marked colored graph, with G connected (and so a single base vertex b). We define $T(G)$ to be $T(\rho(\pi_1(G,b)))$. The proof of the following lemma is similar to that of Lemma 3.3.

Lemma 3.4. *Let (G,γ,B) be a marked colored graph. The quantity $T(G)$ is independent of the choice of base vertices, and so is a property of the underlying colored graph (G,γ).*

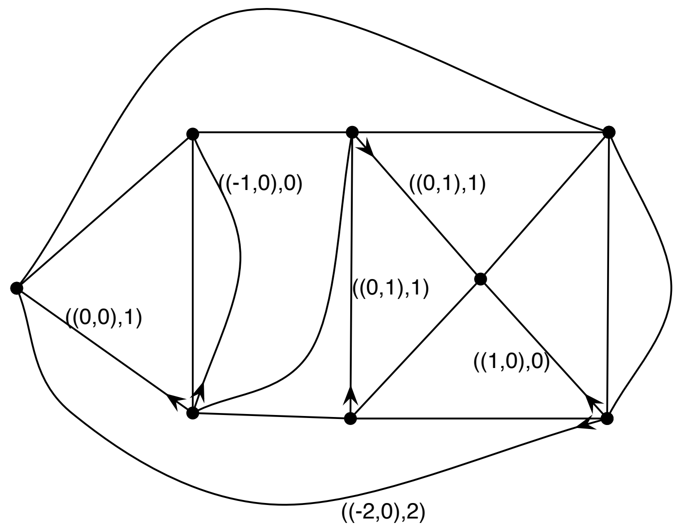

Figure 3: An example of a Γ-$(2,2)$ graph when $\Gamma = \Gamma_3$.

3.3.3. Γ-$(2,2)$ graphs We are now ready to define Γ-$(2,2)$ graphs. Let (G,γ) be a colored graph with n vertices and c connected components G_i. We define the function f to be

$$f(G) = 2n + \mathrm{rep}_{\Gamma_k}(G) - \sum_{i=1}^{c} T(G_i)$$

A colored graph (G,γ) on n vertices and m edges is defined to be a Γ-$(2,2)$ *graph* if:

- The number of edges m is $2n + \mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma_k))$ (i.e., it is the maximum possible value for f).
- For every subgraph G' of G, with m' edges, $m' \le f(G')$.

We note that it is essential that the definition is made over *all* subgraphs, and not just vertex-induced or connected ones. Figure 3 shows an example of a Γ-$(2,2)$ -graph.

3.3.4. Direction network derivation Before continuing with the development of the combinatorial theory, we quickly motivate the definition of Γ-$(2,2)$ -graphs. Readers who are not familiar with rigidity and direction networks may want to either skip to Section 3.3.6 and revisit this, purely informative, section after reading the definitions in Section 4.

Proposition 4.2, in Section 4.1 below, implies that a generic direction network on a Γ_k-colored graph (G,γ) has only *collapsed* realizations (with all the points on top of each other and a trivial representation for $\Lambda(\Gamma_k)$) if and only if (G,γ) has a spanning Γ-$(2,2)$ subgraph.

The definition of the function f comes from analyzing the degrees of freedom in realizations which have the endpoints of each edge coincident (these are collapsed when G is connected). For any realization $G(\mathbf{p},\Phi)$, we can translate it (this preserves directions), so that $\Phi(r_k)$ has the origin as its rotation center. Then, restricted to a subgraph G' of G:

- The total number of variables involved in the equations giving the edge directions is $2n' + \mathrm{rep}_{\Gamma_k}(G')$. Since we fix $\Phi(r_k)$ to rotate at the origin (see Section 4.1 for an explanation

why we can do this), the only variability left in Φ is $\Phi(\Lambda(\Gamma_k))$. Since $\operatorname{rep}_{\Gamma_k}(G')$ measures how much of $\Lambda(\Gamma)$ is "seen" by G', this is the term we add.

- Each connected component G_i has a $T(G_i')$-dimensional space of collapsed realizations. If G_i' has a rotation, then a collapsed realization of the lift $\tilde{G}_i'$ must lie on the corresponding rotation center since a solution must be rotationally symmetric. When G_i' has no rotation, no such restriction exists, and there are 2-dimensions worth of places to put the collapsed $\tilde{G}_i'$. Each collapsed connected component is independent of the others, so this term is additive over connected components.

The heuristic above coincides with the definition of the function f.

3.3.5. Map-graphs In this section we recall the definition of a map-graph. As we will see in the next section, the structure of Γ-$(2,2)$ graphs is closely related to map-graphs. A *map-graph* is a graph in which every connected component has exactly one cycle. In this definition, self-loops correspond to cycles. A 2-*map-graph* is a graph that is the edge-disjoint union of two spanning map-graphs. Observe that map-graphs, and, consequently, 2-map-graphs, do *not* need to be connected.

3.3.6. Γ-$(1,1)$ graphs We will characterize Γ-$(2,2)$ graphs in terms of decompositions into simpler Γ-$(1,1)$ *graphs*[1], which we now define.

Let (G,γ) be a colored graph and select a base vertex b_i for each connected component G_i of G. We define (G,γ) to be a Γ-$(1,1)$ *graph* if:

- G is a map-graph plus $\frac{1}{2}\operatorname{rep}_{\Gamma_k}(\Lambda(\Gamma_k))$ additional edges.
- For each connected component G_i of G, $\rho(\pi_1(G_i, b_i))$ contains a rotation.
- We have $\operatorname{rep}_{\Gamma_k}(G) = \operatorname{rep}_{\Gamma_k}(\Lambda(\Gamma_k))$, or equivalently, $\operatorname{cl}(\Lambda(G,B)) = \Lambda(\Gamma_k)$.

Although we do not define Γ-$(1,1)$ graphs via sparsity counts, there is an alternative characterization in these terms. We define the function $g(G)$ to be

$$g(G) = n + \frac{1}{2}\operatorname{rep}_{\Gamma_k}(G) - \frac{1}{2}\sum_{i=1}^{c} T(G_i)$$

where (G,γ) is a colored graph and n and c are the number of vertices and connected components. Notice that $g = \frac{1}{2}f$. In Section 3.5 we will show:

Proposition 3.5. *The family of Γ-$(1,1)$ graphs gives the bases of a matroid, and the rank of the Γ-$(1,1)$ matroid is given by the function:*

$$g(G) = n + \frac{1}{2}\operatorname{rep}_{\Gamma_k}(G) - \frac{1}{2}\sum_{i=1}^{c} T(G_i)$$

In particular, this implies that g is non-negative, submodular, and monotone.

[1] The terminology of "$(2,2)$" and "$(1,1)$" comes from the fact that spanning trees of finite graphs are "$(1,1)$-tight" in the sense of [10]. The Γ-$(1,1)$ graphs defined here are, in a sense made more precise in [16, Section 5.2], analogous to spanning trees.

3.3.7. Decomposition characterization of Γ-$(2,2)$ graphs The key combinatorial result about Γ-$(2,2)$ graphs, that is used in an essential way to prove the "collapsing lemma" Proposition 4.2, is the following.

Proposition 3.6. *Let (G,γ) be a colored graph. Then (G,γ) is a Γ-$(2,2)$ graph if and only if it is the edge-disjoint union of two spanning Γ-$(1,1)$ graphs.*

Proof. Since $f = 2g$, Proposition 3.5 implies that g meets the hypothesis required for the Edmonds-Rota construction [8], from which we conclude that Γ-$(2,2)$ graphs are a matroidal family. The existence of the desired decomposition follows from the Matroid Union Theorem for rank functions [8]. $\square$

3.4 Γ-colored Laman graphs

We are now ready to define Γ-*colored-Laman graphs*, which are the colored graphs characterizing minimally rigid generic frameworks in Theorem 1. Just as for Γ-$(2,2)$ graphs, we define them via sparsity counts.

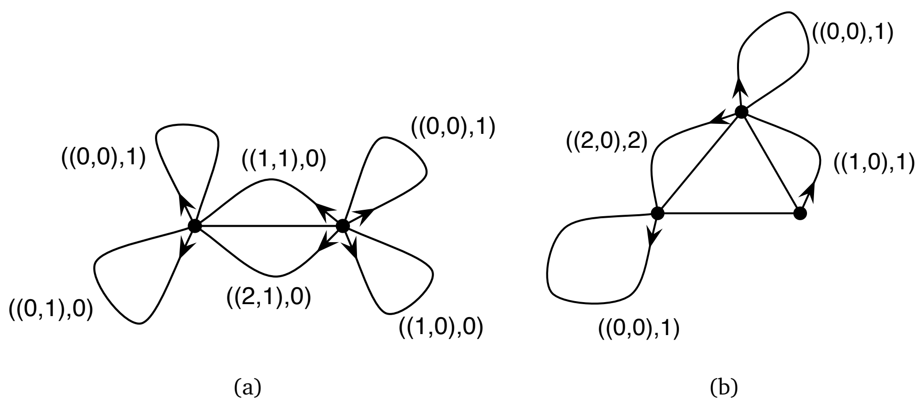

Figure 4: Examples of Γ-colored-Laman graphs: (a) a Γ_2-colored-Laman graph; (b) a Γ_3-colored-Laman graph

3.4.1. Definition of Γ-colored-Laman graphs Let (G,γ) be a colored graph, and let f be the sparsity function defined in Section 3.3. The most direct definition of the sparsity function h for Γ-colored-Laman graphs is:

$$h(G) = f(G) - 1$$

A colored graph (G,γ) is defined to be Γ-*colored-Laman* if:

- G has n vertices and $m = 2n + \mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma_k)) - T(\Gamma_k) - 1$ edges.
- For all subgraphs G' spanning m' edges, $m' \leq h(G')$

Figure 4 shows some examples of Γ-colored-Laman graphs. If a colored graph is a subgraph of a Γ-colored-Laman graph, then it is defined to be Γ-*colored-Laman sparse*. Equivalently, (G,γ) is Γ-colored-Laman sparse when the condition "$m' \leq h(G')$" above holds for all subgraphs G'.

3.4.2. Alternate formulation of Γ-colored-Laman graphs While the definition of h is all that is needed to prove Theorem 1, it does not give any motivation in terms of a degree-of-freedom count. We now give an alternate formulation of Γ-colored-Laman via the Teichmüller space and the centralizer, which were defined in Section 2.5, that will let us do this.

Let (G,γ,B) be a marked colored graph with connected components $G_1,\ldots,G_c$ and n vertices, and let $\Lambda(G,B)$ be its translation subgroup as defined in Section 3.3.1. We define

$$\mathrm{teich}_{\Gamma_k}(G) = \mathrm{teich}_{\Gamma_k}(\Lambda(G,B))$$

which, by a proof nearly identical to that of Lemma 3.3, is well-defined and independent of the choice of base vertices.

For a component G_ℓ with base vertex b_ℓ, we set $\mathrm{cent}_{\Gamma_k}(G_\ell) = \mathrm{cent}_{\Gamma_k}(\rho(\pi_1(G_\ell, b_\ell)))$. For similar reasons, $\mathrm{cent}_{\Gamma_k}(G_\ell)$ is also independent of the base vertex.

We can now define a "more natural" sparsity function

$$h'(G) = 2n + \mathrm{teich}_{\Gamma_k}(G) - \left(\sum_{i=1}^{c} \mathrm{cent}_{\Gamma_k}(G_i))\right)$$

The class of colored graphs defined by h' is the same as that arising from h, giving a second definition of Γ-colored-Laman graphs. Since Lemma 3.7 is not used to prove any further results, we omit the proof.

Lemma 3.7. *A colored graph (G,γ) is Γ-colored-Laman if and only if:*

- *G has n vertices and $m = 2n + \mathrm{teich}(\Gamma) - \mathrm{cent}(\Gamma)$ edges.*
- *For all subgraphs G' spanning m' edges, $m' \le h'(G')$*

3.4.3. Degree of freedom heuristic The function h' is amenable to an interpretation that allows us, by Lemma 3.7, to give a rigidity-theoretic "degree of freedom" derivation of Γ-colored-Laman graphs. This section is expository, and readers unfamiliar with rigidity theory may skip to Section 3.4.4 and return here after reading Section 5.

Given a framework with underlying colored graph (G,γ), with the graph G having n vertices and c connected components $G_1, G_2, \ldots, G_c$, we find that:

- We have $2n$ degrees of freedom from the points. From the representation $\Phi : \Gamma_k \to \mathrm{Euc}(2)$, there are $\mathrm{rep}(\Gamma_k)$ degrees of freedom, but if we mod out by trivial motions from Euc(2), we have $\mathrm{teich}_{\Gamma_k}(\Gamma)$ degrees of freedom left. However, we have only $\mathrm{teich}_{\Gamma_k}(G)$ degrees of freedom that apply to G.
- Each connected component has $\mathrm{cent}_{\Gamma_k}(G_i)$ trivial degrees of freedom. Since elements in the centralizer for G_i commute with those in $\rho(\pi_1(G_i))$, we may "push the vertices of G_i around" with the centralizer elements while preserving symmetry. Since these motions always exist, they are trivial.

This heuristic corresponds to the function h'.

3.4.4. Edge-doubling characterization of Γ-colored-Laman graphs The main combinatorial fact about Γ-colored-Laman graphs we need is the following simple characterization by edge-doubling (cf. [11, 18]).

Proposition 3.8. *Let $\Gamma = \Gamma_k$ for $k = 2, 3, 4$, or 6 be a crystallographic group and let (G, γ) be a Γ-colored graph. Then (G, γ) is Γ-colored-Laman if and only if for any edge $ij \in E(G)$, the colored graph (G', γ') obtained by adding a copy of ij to G with the same color results in a Γ-$(2, 2)$ graph.*

Proof. This is straightforward to check once one notices that (G, γ) is Γ-colored-Laman if and only if no subgraph G' with m' edges has $m' = f(G')$. □

3.4.5. Γ-colored-Laman circuits Let (G, γ) be a colored graph. We define (G, γ) to be a Γ-*colored-Laman circuit* if it is edge-wise minimal with the property of not being Γ-colored-Laman sparse. More formally, (G, γ) is a Γ-colored-Laman circuit if:

- (G, γ) is not Γ-colored-Laman sparse
- For all colored edges $ij \in E(G)$, $(G - ij, \gamma)$ is Γ-colored-Laman sparse

As the terminology suggests, Γ-colored-Laman circuits are the circuits of the matroid that has, as its bases, Γ-colored-Laman graphs. The following lemmas are immediate from the definition.

Lemma 3.9. *Let (G, γ) be a colored graph. If (G, γ) is not Γ-colored-Laman sparse, then it contains a Γ-colored-Laman circuit as a subgraph.*

Lemma 3.10. *Let (G, γ) be a colored graph with n vertices and m edges. Then (G, γ) is a Γ-colored-Laman circuit if and only if:*

- *The number of edges $m = f(G)$*
- *For all subgraphs G' of G, on m' edges, $m' < f(G')$*

Here, f is the colored-$(2, 2)$ sparsity function defined in Section 3.3.

3.5 Γ-$(1, 1)$ graphs: proof of Proposition 3.5

With the definitions and main properties of Γ-$(2, 2)$ and Γ-colored-Laman graphs developed, we prove:

Proposition 3.5. *The family of Γ-$(1, 1)$ graphs gives the bases of a matroid, and the rank of the Γ-$(1, 1)$ matroid is given by the function:*

$$g(G) = n + \frac{1}{2}\operatorname{rep}_{\Gamma_k}(G) - \frac{1}{2}\sum_{i=1}^{c} T(G_i)$$

In particular, this implies that g is non-negative, submodular, and monotone.

Proof. The proposition follows immediately from Lemmas 3.14 and 3.17, which are proven below. □

With this, the proof of Proposition 3.6 is also complete. The rest of this section is organized as follows: first we prove that the Γ-$(1,1)$ graphs give the bases of a matroid and then we argue that the rank function of this matroid is, in fact, the function g, defined in Section 3.3.

We recall from Section 3.2 that, for a marked colored graph (G,γ,B) with a fixed spanning forest F, the map ρ, defined in Section 3.1, induces a map from (G,γ,B,F) to $E_{\Gamma_k,n}$, the ground set of the matroid $M_{\Gamma_k,n}$ from Section 2.6. We adopt the notation of Section 3.2, and denote the image of this map by $A(G,B,F)$.

We start by studying $A(G,B,F)$ in more detail.

3.5.1. Rank of $A(G,B,F)$ As defined, the set $A(G,B,F)$ depends on a choice of base vertices for each connected component and a spanning forest F of G. Since we are interested in constructing a matroid on colored graphs without additional data, the first structural lemma is that the rank of $A(G,B,F)$ in $M_{\Gamma_k,n}$ is independent of the choices for B and F.

Lemma 3.11. *Let (G,γ,B) be a marked colored graph with connected components $G_1,\dots,G_c$ and fix a spanning forest F of G. Then the rank of $A(G,B,F)$ in the matroid $M_{\Gamma_k,n}$ is invariant under an arbitrary change of base vertices and spanning forest.*

Proof. For convenience, shorten the notation $A(G,B,F)$ to A. By Lemma 3.2 $\rho(\pi_1(G_\ell,v_\ell)) = \Gamma_{A,\ell}$. Changing the spanning forest F just picks out a different set of generators for $\pi_1(G_\ell,v_\ell)$, and so does not change $\Gamma_{A,\ell}$, and thus the rank in $M_{\Gamma_k,n}$, which does not depend on the generating set, is unchanged.

To complete the proof, we show that changing the base vertices corresponds, in $E_{\Gamma_k,n}$, to applying the conjugation operation defined in Section 2.6 to A. Suppose that G is connected and fix a spanning tree F and a base vertex b. If P is a closed path starting and ending at b, for any other vertex b' there is a path P' that: starts at b', goes to b along a path $P_{bb'}$, follows P, and then returns from b to b' along $P_{bb'}$ in the other direction. We have $\rho(P') = \rho(P_{bb'})\rho(P)\rho(P_{bb'})^{-1}$, so P and P' have conjugate images. Thus changing base vertices corresponds to conjugation, and by Lemma 2.22 we are done after considering connected components one at a time. □

In light of Lemma 3.11, when we are interested only in the rank of $A(G,B,F)$, we can freely change B and F. Thus, we shorten the notation to $A(G)$.

3.5.2. The effect on $A(G)$ **of adding or deleting a colored edge** In the proof of the basis exchange property, we will need to start with a Γ-$(1,1)$ graph, and add a colored edge to it. There are two possibilities: the edge ij is in the span of some connected component G_i of G or it is not. Each of these has an interpretation in terms of how $A(G+ij)$ is different from $A(G)$.

Lemma 3.12. *Let (G,γ) be a colored graph and let ij be a colored edge. Then:*

- *If the edge ij is in the span of a connected component, G_ℓ of G, then $A(G+ij)$ is $A(G)+(\gamma,\ell)$, where γ is the image of the fundamental closed path of ij with respect to some spanning tree and base vertex of G_ℓ.*
- *If the edge ij connects two connected components G_ℓ and G_r of G, then $A(G+ij)$ is a fusing operation (defined in Section 2.6) on $A(G)$ after a conjugation. In particular, in the notation of Section 2.6, $A(G)_\ell$ and $A(G)_r$ are fused. Conversely, $A(G)$ is a conjugation of a separation of $A(G+ij)$.*

Proof. The first part follows from the fact that if we pick a base vertex and spanning tree of G_ℓ, then adding the colored edge ij to G_ℓ induces exactly one new fundamental closed path.

For the second part, w.l.o.g., assume that G has two connected components G_1, G_2 and that ij connects them. Let T_1 and T_2 be the spanning trees and b_1 and b_2 the base vertices which define the set $A(G)$. Then, we can choose $T = T_1 \cup T_2 + ij$ as the spanning tree for G and b_1 for the base vertex. The fundamental closed paths for edges in $E(G_1) - E(T)$ are unchanged. The ρ-image of the fundamental paths for $E(G_2) - E(T)$ are conjugated by $\rho(P_{b_1 b_2})$ where $P_{b_1 b_2}$ is the unique path in T from b_1 to b_2. Thus $A(G+ij)_1$ consists of $A(G)_1$ and a conjugation of $A(G)_2$. The converse is clear since the inverse of a conjugation is a conjugation, and the inverse of fusing is separating. $\square$

3.5.3. Γ-$(1,1)$ **graphs and tight independent sets in** $M_{\Gamma_k,n}$ Γ-$(1,1)$ graphs (G,γ) have a simple characterization in terms of $A(G)$: they correspond exactly to the situations in which $A(G)$ is tight and independent.

Lemma 3.13. *Let (G,γ) be a colored graph. Then (G,γ) is Γ-$(1,1)$ if and only if $A(G)$ is tight and independent in $M_{\Gamma_k,n}$.*

Proof. We recall that Lemma 2.21 gives a structural characterization of tight independent sets in $M_{\Gamma_k,n}$. The proof proceeds by translating the definitions from Section 2.6.7 into graph theoretic terms. In this proof, we adopt the notation of Section 2.6.7, and we remind the reader that a subset $A \subset E_{\Gamma_k,n}$ is tight if it is independent in $M_{\Gamma_k,n}$ and has

$$|A| = c(A) + \frac{1}{2}\operatorname{rep}_{\Gamma_k}(\Lambda(\Gamma_k))$$

elements.

We first suppose that $A(G)$ is tight, and show that (G,γ) is a Γ-$(1,1)$ graph. By construction $A(G)$ has an element (γ,ℓ) if and only if there is some edge ij in the connected component G_ℓ not in the spanning forest F used to compute $A(G)$. It then follows that, if $A(G)$ is tight, each connected component of G_i of G has at least one more edge than $G_i \cap F$. This implies that G contains a spanning map graph. Because $|A(G)| = c(A) + \frac{1}{2}\operatorname{rep}_{\Gamma_k}(\Lambda(\Gamma_k))$ it then follows that G is a map-graph plus $\frac{1}{2}\operatorname{rep}_{\Gamma_k}(\Lambda(\Gamma_k))$ additional edges, which are the combinatorial hypotheses for being a Γ-$(1,1)$ graph.

Now we use the fact that $A(G)$ is independent in $M_{\Gamma_k,n}$. Independence implies that, if non-empty, $A(G_i)$ contains a rotation, from which it follows that, for each connected component G_i of G, $\rho(\pi_1(G_i,b_i))$ does as well. Similarly, independence of $A(G)$ implies that $\operatorname{rep}_{\Gamma_k}(\Lambda(A(G))) = \operatorname{rep}_{\Gamma_k}(\Lambda(\Gamma_k))$, so $\operatorname{rep}_{\Gamma_k}(G) = \operatorname{rep}_{\Gamma_k}(\Lambda(\Gamma_k))$. We have now shown that (G,γ) is a Γ-$(1,1)$ graph.

The other direction is straightforward to check. $\square$

3.5.4. Γ-$(1,1)$ **graphs form a matroid** We now have the tools to prove that the Γ-$(1,1)$ graphs form the bases of a matroid. We take as the ground set the graph $K_{\Gamma_k,n}$ on n vertices that has one copy of each possible directed edge ij or self-loop ij with color $\gamma \in \Gamma_k$.

Lemma 3.14. *The set of Γ-$(1,1)$ graphs on n vertices form the bases of a matroid on $K_{\Gamma_k,n}$.*

Proof. We check the basis axioms [17, Section 1.2]:
Non-triviality: An uncolored tree plus $\frac{1}{2}\text{rep}_{\Gamma_k}(\Gamma_k)+1$ edges, each of which is colored by a standard generator for Γ_k is clearly Γ-$(1,1)$. Thus the set of bases is not empty.
Equal size: By definition, all Γ-$(1,1)$ graphs have the same number of edges.
Basis exchange: The more difficult step is checking basis exchange. To do this we let G be a Γ-$(1,1)$ graph and ij a colored edge of some other Γ-$(1,1)$ graph which is not in G. It is sufficient to check that there is some colored edge $i'j' \in E(G)$ such that $G+ij-i'j'$ is also a Γ-$(1,1)$ graph. Let (G',γ') be the colored graph $(G+ij,\gamma)$.

Suppose the new edge ij is not a self-loop. Then, pick base vertices B and a spanning forest F of G' that contains the new edge ij. By Lemma 3.11, changing F so as to include ij does not change the rank of $A(G',B,F)$ in $M_{\Gamma_k,n}$. Lemma 3.12 implies that $A(G',B,F)$ is spanning, but not independent, in $M_{\Gamma_k,n}$. Thus there is an element of $A(G',B,F)$ that can be removed to leave a tight, independent set. Since ij is in F, this element does not correspond to ij. The basis exchange axiom then follows from the characterization of Γ-$(1,1)$ graphs in Lemma 3.13.

Suppose ij is a self-loop. Then, the first conclusion of Lemma 3.12 applies. Since ij comes from some other Γ-$(1,1)$ graph, it has non-trivial color, and, thus is not dependent as a singleton set. It follows that there is some element in $A(G')$ (not corresponding to ij) which can be removed to give a tight independent set of the matroid $M_{\Gamma_k,n}$. Consequently, removing the corresponding edge in G' leaves a Γ-$(1,1)$ graph by Lemma 3.13. $\square$

3.5.5. The rank function of the Γ-$(1,1)$ matroid Now we compute the rank function of the Γ-$(1,1)$ matroid. The following lemma is immediate from the definitions.

Lemma 3.15. *Let (G,γ) be a colored graph with n vertices and c connected components. Then*

$$g(G) = n - c + g_1(A(G))$$

where g_1 is the rank function of the matroid $M_{\Gamma_k,n}$.

We can use this to show:

Lemma 3.16. *Let (G,γ) be a colored graph that is independent in the Γ-$(1,1)$ matroid with m edges. Then $m = g(G)$.*

Proof. By definition (G,γ) is a subgraph of some Γ-$(1,1)$ graph (G',γ'). By Lemma 3.13, $m' = g(G')$, where m' is the number of edges of G'. It suffices to show that deleting an edge preserves this equality and independence of $A(G')$. By Lemma 3.12, deleting an edge is equivalent to either removing an element from $A(G')$ or separating and conjugating $A(G')$ and these both preserve independence of $A(G')$. In the first case, $g_1(A(\cdot))$ drops by 1 while n' and c' remain constant, and in the second case n' and $g_1(A(\cdot))$ remain constant while c' increases by 1. $\square$

We can now compute the rank function of the Γ-$(1,1)$ matroid.

Lemma 3.17. *The function g is the rank function of the Γ-$(1,1)$ matroid.*

Proof. Let (G,γ) be an arbitrary colored graph with n vertices and c connected components. The rank of (G,γ) in the Γ-$(1,1)$ matroid is equal to the maximum size of the intersection of G with a Γ-$(1,1)$ graph. Lemma 3.16 implies that what we need to show is that a maximal independent subgraph (G',γ) of (G,γ) has $g(G)$ edges.

We construct G' as follows. First pick a base vertex for every connected component of G and a spanning forest F of G. Initially set G' to be F. Then add edges one at a time to G' from $G - F$ so that $A(G')$ remains independent in $M_{\Gamma_k,n}$ until the rank of $A(G')$ is equal to that of $A(G)$. This is possible by the matroidal property of $M_{\Gamma_k,n}$ and Lemma 3.11, which says the rank of $A(G')$ is invariant under the choices of spanning forest and base vertices.

When the process stops, $A(G')$ is independent in $M_{\Gamma_k,n}$, so G' is independent in the Γ-$(1,1)$ matroid by Lemma 3.13. By construction G' has

$$m' = n - c + g_1(A(G))$$

edges, which is $g(G)$ by Lemma 3.15. □

3.6 Cone-$(1,1)$ and cone-$(2,2)$ graphs

In the proof of Theorem 2 (specifically, Section 4.2 below), we will require some results on direction networks with rotational symmetry from [13, 15]. The combinatorial setup is given in this short section.

3.6.1. Cone-$(1,1)$ graphs Let (G,γ) be a graph whose edges are colored by elements of the group $\mathbb{Z}/k\mathbb{Z}$. As before, there is a well-defined map $\rho : \pi_1(G_i, b_i) \to \mathbb{Z}/k\mathbb{Z}$ where b_i is a vertex in the connected component G_i of G. We define (G,γ) to be a cone-$(1,1)$ graph if G is a map-graph and the cycle in each connected component has non-trivial ρ-image. We define the quantity $T(G_i)$ to be the same one defined in Section 3.3, where all nontrivial elements of $\mathbb{Z}/k\mathbb{Z}$ are "rotations".

The sparsity characterization of cone-$(1,1)$ graphs is:

Lemma 3.18 ([15, Section 2.6], [27, "Matroid Theorem"]). *The cone-$(1,1)$ graphs on n vertices are the bases of a matroid that has as its rank function*

$$r(G') = n' - \frac{1}{2}\sum_{i=1}^{c} T(G_i)$$

where n' and c' are the number of vertices and connected components in G'.

3.6.2. Cone-$(2,2)$ graphs Let (G,γ) be a $\mathbb{Z}/k\mathbb{Z}$ colored graph with n vertices. We define (G,γ) to be a cone-$(2,2)$ graph if:

- G has $m = 2n$ edges.
- For all subgraphs with m' edges, n' vertices, and connected components $G_1, G_2, \ldots, G_c$,

$$m' \leq 2n' - \sum_{i=1}^{c} T(G_i)$$

If only the second condition holds, then (G,γ) is defined to be *cone-$(2,2)$ sparse*.

3.7 Generalized cone-$(2,2)$ graphs

As a technical tool in the proof of Theorem 2, we will use *generalized cone-$(2,2)$ graphs*. These are Γ_k-colored graphs, which we will define in terms of a decomposition property.

3.7.1. Generalized cone-$(1,1)$ graphs Let (G,γ) be a Γ_k-colored graph. We define (G,γ) to be a *generalized cone-$(1,1)$ graph* if, after considering the ρ-image modulo the translation subgroup, the result is a cone-$(1,1)$ graph. Equivalently, (G,γ) is a generalized cone-$(1,1)$ graph if:

- G is a map graph
- The ρ-image of the cycle in each connected component of G is a rotation

The difference between cone-$(1,1)$ graphs and generalized cone-$(1,1)$ graphs is that the rotations need not be around the same center. By modding colors out by $\Lambda(\Gamma_k)$, the next lemma follows easily from Lemma 3.18.

Lemma 3.19. *The generalized cone-$(1,1)$ graphs on n vertices are the bases of a matroid that has as its rank function*

$$r(G') = n' - \frac{1}{2}\sum_{i=1}^{c} T(G_i)$$

where n' and c' are the number of vertices and connected components in G'.

3.7.2. Relation to Γ-$(1,1)$ graphs Generalized cone-$(1,1)$ graphs are related to Γ-$(1,1)$ graphs by this next sequence of lemmas.

Lemma 3.20. *Let (G,γ) be a Γ-$(1,1)$ graph. Then (G,γ) contains a generalized cone-$(1,1)$ graph as a spanning subgraph.*

Proof. This follows from the definition, since each connected component G_i of G has $T(G_i) = 0$. It follows that G_i has a spanning subgraph that is a connected map-graph with its cycle having a rotation as its ρ-image. □

Let (G,γ) be a Γ-$(1,1)$ graph, and let (G',γ) be a spanning generalized cone-$(1,1)$ subgraph. One exists by Lemma 3.20. We define (G',γ) to be a *g.c.-$(1,1)$ basis* of (G,γ).

Lemma 3.21. *Let (G,γ) be a Γ_k-colored Γ-$(1,1)$ graph for $k = 3,4,6$. Let (G',γ) be a g.c.-$(1,1)$ basis of (G,γ), and let ij be the (unique) edge in $E(G) - E(G')$. Then either:*

- *The colored edge ij is a self-loop and the color γ_{ij} is a translation.*
- *There is a unique minimal subgraph G'' of G, such that the ρ-image of (G'',γ) includes a translation, ij is an edge of G'', and if $vw \in E(G'')$, then $(G' + ij - vw,\gamma)$ is also a g.c.-$(1,1)$ basis of (G,γ).*

Proof. If ij is a self-loop colored by a translation, then it is a circuit in the matroid of generalized cone-$(1,1)$ graphs on the ground set (G,γ). Otherwise, the subgraph G'' which the lemma requires is just the fundamental generalized cone-$(1,1)$ circuit of ij in (G',γ). □

3.7.3. Generalized cone-$(2,2)$ graphs Let (G,γ) be a Γ_k-colored graph. We define (G,γ) to be a *generalized cone-*$(2,2)$ *graph* if it is the union of two generalized cone-$(1,1)$ graphs. Using the Edmonds-Rota construction [8], the same way we did in Section 3.3.7, we get:

Lemma 3.22. *The generalized cone-*$(2,2)$ *graphs on* n *vertices give the bases of a matroid.*

The other fact about generalized cone-$(2,2)$ graphs is their relationship to Γ-$(2,2)$ graphs.

Lemma 3.23. *Let* (G,γ) *be a* Γ-$(2,2)$ *graph. Then* (G,γ) *contains a generalized cone-*$(2,2)$ *graph as a spanning subgraph.*

Proof. This follows from Proposition 3.6 and Lemma 3.20. □

4. Direction networks

4.1 Crystallographic direction networks

Let $(\tilde{G},\varphi)$ be a graph with a Γ_k-action φ. A *crystallographic direction network* $(\tilde{G},\varphi,\tilde{\mathbf{d}})$ is given by $(\tilde{G},\varphi)$ and an assignment of a direction $\tilde{\mathbf{d}}_{ij} \in \mathbb{R}^2 \setminus \{0\}$ to each edge $ij \in E(\tilde{G})$.

We will, moreover, require that the direction networks themselves be symmetric in the following sense. For any $\Phi \in \operatorname{Rep}(\Gamma)$ and any $\gamma \in \Gamma$, the rotational part $\Phi(\gamma)_r$ of $\Phi(\gamma)$ depends only on γ and not Φ. Thus, we will require the directions to be equivariant with respect to this action; i.e., if $i'j' = \gamma \cdot ij$, then $\mathbf{d}_{i'j'} = \Phi(\gamma)_r \mathbf{d}_{ij}$.

A direction network on $(\tilde{G},\varphi)$, thus, is completely determined after assigning a direction to one edge in each Γ-orbit. It is plain, then, that this is equivalent to assigning directions to the colored quotient graph (G,γ). The dictionary is straightforward as well. Given $(\tilde{G},\varphi,\mathbf{d})$, the direction for edge ij of G is the same as the direction of the chosen edge representative in its fiber.

4.1.1. The realization problem A *realization* of a crystallographic direction network is given by a point set $\mathbf{p} = (\mathbf{p}_i)_{i\in V(\tilde{G})}$ and a representation $\Phi \in \overline{\operatorname{Rep}}(\Gamma_k)$:

$$\left\langle \mathbf{p}_j - \mathbf{p}_i, \tilde{\mathbf{d}}_{ij}^{\perp} \right\rangle = 0 \quad \text{for all edges } ij \in E(\tilde{G}) \tag{7}$$

$$\mathbf{p}_{\gamma\cdot i} = \Phi(\gamma)\cdot \mathbf{p}_i \quad \text{for all vertices } i \in V(\tilde{G}) \tag{8}$$

We denote realizations by $\tilde{G}(\mathbf{p},\Phi)$, to indicate the dependence on Φ.

We define now *collapsed* and *faithful* realizations. An edge ij is *collapsed* in a realization $\tilde{G}(\mathbf{p},\Phi)$ if $\mathbf{p}_i = \mathbf{p}_j$. A *realization is collapsed* when all the edges are collapsed *and* Φ is trivial. A representation is *trivial* if it maps $\Lambda(\Gamma_k)$ to zero. A realization is *faithful* if no edge is collapsed and Φ is not trivial.

4.1.2. Direction Network Theorem Our main theorem on crystallographic direction networks is the following.

Theorem 2. *Let* Γ *be an orientation-preserving crystallographic group. A generic* Γ*-crystallographic direction network* $(\tilde{G},\varphi,\tilde{\mathbf{d}})$ *has a unique, up to translation and scaling, faithful realization if and only if its associated colored graph is* Γ*-colored-Laman.*

For technical simplification, we reduce Theorem 2 to the following proposition which is the same result except that the rotation center of $\Phi(r_k)$ is fixed to be the origin.

Proposition 4.1. *Let Γ be an orientation-preserving crystallographic group. A generic Γ-crystallographic direction network $(\tilde{G}, \varphi, \tilde{\mathbf{d}})$ has a unique, up to scaling, faithful realization $\tilde{G}(\mathbf{p}, \Phi)$ satisfying $\Phi(r_k) = (0, R_k)$ if and only if its associated colored graph is Γ-colored-Laman.*

Theorem 2 follows easily from the proposition.

Proof of Theorem 2 from Proposition 4.1. Let $(\tilde{G}, \varphi, \tilde{\mathbf{d}})$ be any generic direction network. For any realization $\tilde{G}(\mathbf{p}, \Phi)$ of $(\tilde{G}, \varphi, \tilde{\mathbf{d}})$ and any translation ψ, we have that $\tilde{G}(\psi(\mathbf{p}), \Phi^{\psi})$ is also realization of $(\tilde{G}, \varphi, \tilde{\mathbf{d}})$ where Φ^{ψ} is the representation defined by $\Phi^{\psi}(\gamma) = \psi\Phi(\gamma)\psi^{-1}$. In particular, for any realization $\tilde{G}(\mathbf{p}, \Phi)$, there is a unique translation ψ such that $\psi\Phi(r_k)\psi^{-1} = (0, R_k)$. $\square$

Our goal then is to prove Proposition 4.1. For the remainder of this section, we will require all realizations to map the rotational generator r_k of Γ_k to the counter-clockwise rotation around the origin R_k of angle $2\pi/k$. All of the propositions in the remainder of this section operate under this assumption. However, note that some of the results about direction networks, including Proposition 4.3, hold without this restriction.

4.1.3. Proof of Proposition 4.1 Let (G, γ) be a colored graph. The key proposition, which is proved in Section 4.3 is the following.

Proposition 4.2. *A generic crystallographic direction network which has a Γ-$(2,2)$ colored quotient graph has only collapsed realizations.*

By counting the dimension of the space of collapsed realizations, done in Section 4.4 below, it follows that:

Proposition 4.3. *A generic crystallographic direction network which has a Γ-colored-Laman circuit as its quotient graph has only realizations with collapsed edges.*

Proposition 4.3 readily implies one direction of Proposition 4.1. If (G, γ) is not a Γ-colored-Laman graph, then it has either too few edges or contains a Γ-colored-Laman circuit as a subgraph. In the former case, a dimension count implies that a faithful realization cannot be unique up to translation and scale and in the latter, every realization contains collapsed edges by Proposition 4.3.

For the other direction, we assume that (G, γ) is Γ-colored-Laman with n vertices. Since every Γ-colored-Laman graph is Γ-$(2,2)$ sparse by Proposition 3.8, we see from Proposition 4.2 that the equations defining the realization space of a generic colored direction network with quotient graph $G(\gamma)$ is 1-dimensional. This means every realization is a rescaling of a single realization, so if any edge ij is collapsed, it is collapsed in all realizations. In particular, if we double ij and assign it a generic direction, the realization space will not change. Proposition 3.8 tells us that the graph obtained in this way is Γ-$(2,2)$, so Proposition 4.2 applies to it, showing that the realization space is zero-dimensional. The resulting contradiction completes the proof. $\square$

4.1.4. Colored direction networks We will make use of *colored crystallographic direction networks* to study crystallographic direction networks. Since there is no chance of confusion, we simply call these "colored direction networks" in the next several sections. A *colored direction network* $(G, \gamma, \mathbf{d})$ is given by a Γ_k-colored graph (G, γ) and an assignment of a direction $\mathbf{d}_{ij}$ to every edge ij. The realization system for $(G, \gamma, \mathbf{d})$ is given by

$$\left\langle \Phi(\gamma_{ij}) \cdot \mathbf{p}_j - \mathbf{p}_i, \mathbf{d}_{ij}^{\perp} \right\rangle = 0 \tag{9}$$

The unknowns are the representation Φ of Γ_k and the points $\mathbf{p}_i$. (As above $\Phi(r_k)$ is restricted to be rotation about the origin.) We denote points in the realization space by $G(\mathbf{p}, \Phi)$. We observe here that, for Φ parameterized by the vectors in Lemma 2.3, the realization system is linear. This can be seen by, e.g., the computations in Section 4.2. The following two lemmas linking crystallographic direction networks and colored direction networks follow easily from Lemma 3.1 and the observations above.

Lemma 4.4. *Given a colored direction network $(G, \gamma, \mathbf{d})$, its lift to a crystallographic direction network $(\tilde{G}, \varphi, \tilde{\mathbf{d}})$ is well-defined and the realization spaces of $(G, \gamma, \mathbf{d})$ and $(\tilde{G}, \varphi, \tilde{\mathbf{d}})$ are isomorphic. In particular, they have the same dimension.*

Lemma 4.5. *Let (G, γ) be a Γ_k-colored graph and $(\tilde{G}, \varphi)$ its lift. Assigning a direction to one representative of each edge orbit under φ in $\tilde{G}$ gives a well defined colored direction network $(G, \gamma, \mathbf{d})$.*

This next lemma, which is also immediate from the definitions, describes collapsed edges in terms of colored direction networks.

Lemma 4.6. *Let $(G, \gamma, \mathbf{d})$ be a colored direction network and let $G(\mathbf{p}, \Phi)$ be a realization of $(G, \gamma, \mathbf{d})$. Let $(\tilde{G}, \varphi, \tilde{\mathbf{d}})$ be the lift of $(G, \gamma, \mathbf{d})$ and $\tilde{G}(\mathbf{p}, \Phi)$ be the associated lift of $G(\mathbf{p}, \Phi)$. Then a colored edge $ij \in E(G)$ lifts to an orbit of collapsed edges in $\tilde{G}(\mathbf{p}, \Phi)$ if and only if*

$$\mathbf{p}_i = \Phi(\gamma_{ij}) \cdot \mathbf{p}_j$$

in $G(\mathbf{p}, \Phi)$.

In light of Lemmas 4.4–4.6, we may switch freely between the formalisms, and we do so in subsequent sections.

4.1.5. A result on cone direction networks We prove Proposition 4.2 by bootstrapping results for generalized cone-$(2,2)$ graphs (defined in Section 3.7). The steps are:

- We show that, for fixed Φ, a generic direction network on a generalized cone-$(2,2)$ graph has a unique solution (Proposition 4.7).
- Then we allow Φ to flex. We show that by adding $\mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma_k))$ edges that extend a generalized cone-Laman graph to a Γ-$(2,2)$ graph, realizations of a generic direction network are forced to collapse.

This is done in Sections 4.2 and 4.3. In the first step, we will make use of a result on finite direction networks with rotational symmetry from [15]. A *cone direction network* $(G, \gamma, \mathbf{d})$ is an

assignment $\mathbf{d}_{ij}$ to each edge ij. A realization $G(\mathbf{p})$ of the direction network is a selection of points $\mathbf{p}_i \in \mathbb{R}^2$ such that

$$\langle R_k^{\gamma_{ij}} \mathbf{p}_j - \mathbf{p}_i, \mathbf{d}_{ij} \rangle = 0 \text{ for all } ij \tag{10}$$

Here R_k is the rotation about the origin that rotates through angle $2\pi/k$. A straightforward application of [15, Proposition 3.1] yields:

Theorem 3 ([15])**.** *The system* (10) *defining a generic cone direction network* $(G, \gamma, \mathbf{d})$ *is independent if and only if* (G, γ) *is cone-*$(2,2)$ *sparse.*

4.2 Direction networks on generalized cone-$(2,2)$ graphs

Let (G, γ) be a generalized cone-$(2,2)$ graph. In light of Theorem 3, it should be unsurprising that the realization system (9) has generic rank $2n$ for a colored direction network on (G, γ), since cone direction networks are a "special case". Here is the precise reduction.

Proposition 4.7. *Fix a representation* Φ *of* Γ_k*. Holding* Φ *fixed, if a generic crystallographic direction network has a generalized cone-*$(2,2)$ *colored quotient, then it has a unique realization.*

Proposition 4.7 is immediate from the following statement and Lemma 4.4.

Proposition 4.8. *Let* (G, γ) *be a generalized cone-*$(2,2)$ *graph with* n *vertices. Then the generic rank of the realization system* (9) *is* $2n$*.*

Proof. Expanding (9) we get

$$\left\langle \Phi(\gamma_{ij}) \cdot \mathbf{p}_j - \mathbf{p}_i, \mathbf{d}_{ij}^{\perp} \right\rangle = \left\langle \Phi(\gamma_{ij}) \cdot \mathbf{p}_j, \mathbf{d}_{ij}^{\perp} \right\rangle - \left\langle \mathbf{p}_i, \mathbf{d}_{ij}^{\perp} \right\rangle \tag{11}$$

Define $\Phi(\gamma_{ij})_r \in SO(2)$ to be the rotational part of $\Phi(\gamma_{ij})$ and $\Phi(\gamma_{ij})_t \in \mathbb{R}^2$ to be the translational part, so that $\Phi(\gamma_{ij}) \cdot \mathbf{p} = \Phi(\gamma_{ij})_r \cdot \mathbf{p} + \Phi(\gamma_{ij})_t$. In this notation, (11) becomes

$$\left\langle \Phi(\gamma_{ij})_r \cdot \mathbf{p}_j, \mathbf{d}_{ij}^{\perp} \right\rangle + \left\langle \Phi(\gamma_{ij})_t, \mathbf{d}_{ij}^{\perp} \right\rangle - \left\langle \mathbf{p}_i, \mathbf{d}_{ij}^{\perp} \right\rangle = 0 \tag{12}$$

Since the rotational part $\Phi(\gamma_{ij})_r$ preserves the inner product, we see that (9) is equivalent to the inhomogeneous system

$$\left\langle \mathbf{p}_j, \Phi(\gamma_{ij}^{-1})_r \cdot \mathbf{d}_{ij}^{\perp} \right\rangle - \left\langle \mathbf{p}_i, \mathbf{d}_{ij}^{\perp} \right\rangle = - \left\langle \Phi(\gamma_{ij})_t, \mathbf{d}_{ij}^{\perp} \right\rangle \tag{13}$$

The l.h.s. of (13) is equivalent to (10), and thus the generic rank of (13) is at least as large as that of (10). The proposition then follows from Theorem 3. $\square$

4.3 Proof of Proposition 4.2

We now have the tools in place to prove:

Proposition 4.2. *A generic crystallographic direction network which has a* Γ-$(2,2)$ *colored quotient graph has only collapsed realizations.*

The proof is split into two cases, Γ_2 and Γ_k for $k = 3, 4$, or 6.

4.3.1. Proof for rotations of order 3, 4, **or** 6 Let (G,γ) be a Γ-$(2,2)$ graph. We construct a direction network on (G,γ) that has only collapsed solutions, from which the desired generic statement follows.

Assigning directions We select directions $\mathbf{d}$ for each edge in G with coordinates that are algebraically independent over $\mathbb{Q}$; i.e., their coordinates satisfy no polynomial with integer coefficients.

The realization space of any spanning g.c.-$(2,2)$ **basis** With these direction assignments, we can compute the dimension of the realization space for the direction network induced on any spanning g.c.-$(2,2)$ basis of (G,γ) where by g.c.-$(2,2)$ basis we mean a basis in the generalized cone-$(2,2)$ matroid. One exists by Lemma 3.23.

Lemma 4.9. *Let (G',γ) be a spanning g.c.-$(2,2)$ basis of (G,γ). Then the realization space of the induced direction network $(G',\gamma,\mathbf{d})$ is 2-dimensional, and linearly depends on the representation Φ.*

Proof. The dimension comes from Proposition 4.7 and comparing the number of variables to the number of equations in the realization system (9). Moving the variables associated with Φ to the right as in equation (13) completes the proof. □

A g.c.-$(2,2)$ **basis with non-collapsed complement** By edge counts, there are exactly two edges ij and vw in the complement of any g.c.-$(2,2)$ basis of (G,γ). We will show that there are two edges which do not collapse (and more) when enforcing the directions on the complement.

Lemma 4.10. *There are edges ij, vw of G such that $G' = G - ij - vw$ is a g.c.-$(2,2)$ basis and for all vectors $\mathbf{u} \in \mathbb{R}^2$ there is a realisation $G(\mathbf{p},\Phi)$ of $(G',\gamma,\mathbf{d})$ such that $\Phi(\gamma_{ij})\mathbf{p}_j - \mathbf{p}_i = \mathbf{u}$ (similarly there is a realization such that $\Phi(\gamma_{vw})\mathbf{p}_w - \mathbf{p}_v = \mathbf{u}$).*

Proof. By Proposition 3.6, we can decompose (G,γ) into two spanning Γ-$(1,1)$ graphs X and Y. Since Γ-$(1,1)$ graphs are g.c.-$(1,1)$ graphs plus an edge, there are unique g.c.-$(1,1)$ circuits X'' and Y'' in X and Y respectively. Let $v'w'$ be some edge in Y'' and let $Y' = Y - v'w'$.

Suppose, for contradiction, that for an arbitrary edge ij of X'', the vector $\Phi(\gamma_{ij})\mathbf{p}_j - \mathbf{p}_i$ is constrained to a one-dimensional subspace (or smaller) over all realizations $G(\mathbf{p},\Phi)$ of the direction network $(X - ij \cup Y',\gamma,\mathbf{d})$. Then, either the vector is identically zero or by genericity of $\mathbf{d}$, the vector $\mathbf{d}_{ij}$ differs from $\Phi(\gamma_{ij})\mathbf{p}_j - \mathbf{p}_i$. In either case the edge ij is collapsed in all realizations of $(X \cup Y',\gamma,\mathbf{d})$. Since ij was arbitrary in X'', all edges in X'' are collapsed in all realizations of $(X \cup Y',\gamma,\mathbf{d})$.

However, if every edge in X'' is collapsed in every realization of the direction network $(X \cup Y',\gamma,\mathbf{d})$, this implies that Φ must always be trivial in any realization. Proposition 4.7 would then imply that the realization space is 0-dimensional, and this contradicts the fact that it is at least 1-dimensional, by Lemma 4.9. Thus, it must be that for some edge ij in X'', the vector $\Phi(\gamma_{ij})\mathbf{p}_j - \mathbf{p}_i$ sweeps out all of $\mathbb{R}^2$ as $G(\mathbf{p},\Phi)$ varies over all realizations of $(X - ij \cup Y',\gamma,\mathbf{d})$: any direction is achievable by changing Φ and we can scale. Let now $X' = X - ij$.

By reversing the roles of X and Y we can find an edge vw of Y with the same properties. (Note that the ij we chose was in X'' so the situation is symmetric!) □

The representation Φ must be trivial The rest of the proof will be to show that, adding back ij and vw forces all realizations of $(G, \gamma, \mathbf{d})$ to collapse. Lemma 4.10 and Lemma 4.9 tell us that for realizations $G(\mathbf{p}, \Phi)$ of $(G', \gamma, \mathbf{d})$, both vectors $\Phi(\gamma_{ij})\mathbf{p}_j - \mathbf{p}_i$ and $\Phi(\gamma_{vw})\mathbf{p}_w - \mathbf{p}_v$ depend linearly in a one-to-one fashion on Φ which parameterizes the realization space.

Consequently, if we add the edge ij to G', the new direction network must constrain Φ, and thus $\Phi(\gamma_{vw})\mathbf{p}_w - \mathbf{p}_v$, to some one-dimensional space. Since $\mathbf{d}$ was chosen generically, $\mathbf{d}_{vw}$ differs from this latter vector, and thus Φ is trivial in a realization of $(G, \gamma, \mathbf{d})$. Since Φ is trivial, then by Proposition 4.7 the unique realization must be the completely collapsed one. $\square$

4.3.2. Proof for rotations of order 2 Let (G, γ) be a Γ-$(2,2)$ graph. Again, we will assign directions so that the resulting direction network $(G, \gamma, \mathbf{d})$ has only collapsed solutions. The proof has a slightly different structure from the $k = 3, 4, 6$ case. The main geometric lemma is the following.

Lemma 4.11. *Let (X, γ) be a Γ-$(1,1)$ graph with Γ_2 colors, and let $(X, \gamma, \mathbf{d})$ be a colored direction network which assigns the same direction $\mathbf{v}$ to every edge. Then, any realization $X(\mathbf{p}, \Phi)$ lifts to a realization $\tilde{X}(\mathbf{p}, \Phi)$ such that every vertex lies on a single line in the direction of $\mathbf{v}$.*

Proof that Lemma 4.11 implies Proposition 4.2 for Γ_2 With Lemma 4.11, the Proposition follows readily: the combinatorial Proposition 3.6 says we may decompose (G, γ) into two spanning Γ-$(1,1)$ graphs, which we define to be X and Y. We assign the edges of X a direction $\mathbf{v}_X$ and the edges of Y a linearly independent direction $\mathbf{v}_Y$. Applying Lemma 4.11, to X and Y separately shows that every vertex of a lifted realization $\tilde{G}(\mathbf{p}, \Phi)$ must lie in two skew lines. This is possible only when they are all at the intersection of these lines, implying only collapsed realizations. $\square$

Proof of Lemma 4.11 Let (X, γ) be a Γ-$(1,1)$ graph with Γ_2 colors, and let $(X, \gamma, \mathbf{d})$ be a direction network that assigns all the edges the same direction. Let (X', γ) be a spanning g.c.-$(1,1)$ basis of (X, γ); one exists by Lemma 3.20.

Lemma 4.12. *Let $(X', \gamma, \mathbf{d})$ be a g.c.-$(1,1)$ graph, and let $\mathbf{d}$ assign the same direction $\mathbf{v}$ to every edge. Then, in any realization of the lifted crystallographic direction network $(\tilde{X}, \varphi, \mathbf{d})$, every vertex and every edge lies on a line in the direction $\mathbf{v}$ through a rotation center of $\Phi(\gamma)$ for some $\gamma \in \Gamma_k$.*

Proof: It will suffice to prove the lemma when X' is connected. Because the ρ-image of X' contains an order 2 rotation r, for some vertex $i \in V(X')$, there is a vertex $\tilde{i}$ in the fiber over i such that $\mathbf{p}_{\tilde{i}} - \mathbf{p}_{r \cdot \tilde{i}} = \mathbf{p}_{\tilde{i}} - \Phi(r) \cdot \mathbf{p}_{\tilde{i}}$ is in the direction $\mathbf{v}$. Because $\Phi(r)$ is a rotation through angle π, this means that $\mathbf{p}_{\tilde{i}}$ and $\mathbf{p}_{r \cdot \tilde{i}}$ lie on a line through the rotation center of $\Phi(r)$ in the direction $\mathbf{v}$. Because X' is connected, and edge directions (up to sign) are fixed under an order 2 rotation, the same is then true for every vertex in the connected component $\tilde{X}'_0$ of the lifted realization $\tilde{X}(\mathbf{p}, \Phi)$ that contains $\mathbf{p}_{\tilde{i}}$. The lemma then follows by considering translates of $\tilde{X}'_0$. $\square$Lemma 4.12

To complete the proof, we recall that the ρ-image of (X, γ) contains two linearly independent translations t and t'. This implies that in the lifted realization $\tilde{X}(\mathbf{p}, \Phi)$, there is a vertex $\tilde{i}$ connected by a path of edges in $\tilde{X}$ to $t(\tilde{i})$. Since all edges have the same direction in a realization, there is some $\lambda \in \mathbb{R}$ such that $\lambda \mathbf{v} = \mathbf{p}_{t(\tilde{i})} - \mathbf{p}_{\tilde{i}} = \Phi(t) \cdot \mathbf{p}_{\tilde{i}} - \mathbf{p}_{\tilde{i}}$. Thus, $\Phi(t)$ is a translation in the direction of $\mathbf{v}$. The same argument applies to $\Phi(t')$.

From this, it follows that the rotation center of all rotations $\Phi(r)$ must lie on single line. Lemma 4.12 then applies, so we are done. □

4.4 Proof of Proposition 4.3

We now prove the proposition required for the "Maxwell direction" of Theorem 2:

Proposition 4.3. *A generic crystallographic direction network which has a Γ-colored-Laman circuit as its quotient graph has only realizations with collapsed edges.*

In the proof, we will use the following statement. (cf. [16, Lemma 14.2] for the case when the ρ-image is a translation subgroup and $\Gamma = \Gamma_1$.)

Lemma 4.13. *Let $(G, \gamma, \mathbf{d})$ be a colored direction network on a colored graph (G, γ) with connected components $G_1, G_2, \ldots, G_c$. Then $(G, \gamma, \mathbf{d})$ has at least*

$$\mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma_k)) - \mathrm{rep}_{\Gamma_k}(G) + \sum_{i=1}^{c} T(G_i)$$

dimensions of collapsed realizations.

We defer the proof of Lemma 4.13 to Section 4.4.1 and first show how Lemma 4.13 implies Proposition 4.3. Let (G, γ) be a Γ-colored-Laman circuit with n vertices, m edges,and c connected components $G_1, G_2, \ldots G_c$. By Lemma 3.10, we have

$$m = 2n + \mathrm{rep}_{\Gamma_k}(G) - \sum_{i=1}^{c} T(G_i)$$

It follows from Proposition 4.2 that for generic directions, a colored direction network $(G, \gamma, \mathbf{d})$ has a space of realizations with dimension

$$2n + \mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma_k)) - m = \mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma_k)) - \mathrm{rep}_{\Gamma_k}(G) + \sum_{i=1}^{c} T(G_i).$$

Applying Lemma 4.13 shows that in all of them every edge is collapsed. □

4.4.1. Proof of Lemma 4.13 For now, assume that the colored graph (G, γ) is connected. Select a base vertex b.

Representations that are trivial on $\Lambda(G, b)$ Consider $\Phi \in \overline{\mathrm{Rep}_{\Gamma_k}}(\Lambda(\Gamma_k))$ such that

$$\Phi(t) = ((0,0), \mathrm{Id})$$

for all translations $t \in \Lambda(G, b)$. These representations form a $(\mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma_k)) - \mathrm{rep}_{\Gamma_k}(G))$-dimensional space.

Collapsed realizations for a fixed representation Now we show that there are $T(G)$ dimensions of realizations with all edges collapsed. We do this with an explicit construction. There are two cases.

Case 1: $T(G) = 2$**.** In this case, we know that the subgroup generated by $\rho(\pi_1(G,b))$ is a translation subgroup. Fix a spanning tree T of G and a point $\mathbf{p}_b \in \mathbb{R}^2$. We will construct a realization with vertex b mapped to $\mathbf{p}_b$ and all edges collapsed.

For any pair of vertices i and j, define Q_{ij} to be the path in T from i to j and define η_{ij} to be $\rho(Q_{ij})$. We then set $\mathbf{p}_i = \Phi(\eta_{bi}^{-1}) \cdot \mathbf{p}_b$ for all vertices $i \in V(G)$ other than b. Thus all vertex locations are determined by $\mathbf{p}_b$, giving a 2-dimensional space of realizations for this Φ. We need to check that all edges are collapsed. If ij is an edge of T with color γ_{ij}, then we have

$$\gamma_{ij}^{-1} = \eta_{bj}^{-1} \cdot \eta_{bi}$$

Using this relation, we see that

$$\mathbf{p}_j = \Phi(\eta_{bj}^{-1}) \cdot \mathbf{p}_b = \Phi(\gamma_{ij}^{-1} \cdot \eta_{bi}^{-1}) \cdot \mathbf{p}_b = \Phi(\gamma_{ij}^{-1}) \cdot \mathbf{p}_i$$

so the edge ij is collapsed. If ij is not an edge in T, then the fundamental closed path P_{ij} of ij relative to T and b follows Q_{bi}, crosses ij, and returns to b along Q_{jb}. This gives us the relation

$$\gamma_{ij} = \eta_{bi}^{-1} \cdot \rho(P_{ij}) \cdot \eta_{bj}$$

We then compute

$$\Phi(\gamma_{ij}) \cdot \mathbf{p}_j = (\Phi(\eta_{bi}^{-1}) \cdot \Phi(\rho(P_{ij})) \cdot \Phi(\eta_{bj})) \cdot \mathbf{p}_j$$

Since Φ is trivial on the ρ-images of fundamental closed paths, the r.h.s. simplifies to

$$\Phi(\eta_{bi}^{-1}) \cdot \Phi(\eta_{bj}) \cdot \mathbf{p}_j = \Phi(\eta_{bi}^{-1}) \cdot \mathbf{p}_b = \mathbf{p}_i$$

and we have shown that all edges are collapsed.

Case 2: $T(G) = 0$**.** We adopt the notation from Case 1. As before, we fix a spanning tree T and a representation Φ that is trivial on the translation subgroup $\Lambda(G,b)$. By Lemma 2.5, $\rho(\pi_1(G,b))$ is generated by a translation subgroup $\Gamma' < \Lambda(G,b)$ and a rotation $r \in \Gamma_k$. We set $\mathbf{p}_b$ to be on the rotation center of $\Phi(r)$ and define the rest of the $\mathbf{p}_i$ as before: $\mathbf{p}_i = \Phi(\eta_{bi}^{-1}) \cdot \mathbf{p}_b$. Observe that $\Phi(r)$ then fixes $\mathbf{p}_b$.

For edges ij in the tree T, the argument that ij is collapsed from Case 1 applies verbatim. For non-tree edges ij, a similar argument relating the fundamental closed path P_{ij} to Q_{bi} and Q_{bj} yields the relation

$$\gamma_{ij} = \eta_{bi}^{-1} \cdot \rho(P_{ij}) \cdot \eta_{bj}$$

Since Φ is trivial on translations $t \in \Gamma'$, we see that for some ℓ

$$\Phi(\gamma_{ij}) = \Phi(\eta_{bi}^{-1}) \cdot \Phi(r^\ell) \cdot \Phi(\eta_{bj})$$

We then compute

$$\Phi(\gamma_{ij})\mathbf{p}_j = \Phi(\eta_{bi}^{-1}) \cdot \Phi(r^\ell) \cdot \Phi(\eta_{bj}) \cdot \mathbf{p}_j = \Phi(\eta_{bi}^{-1}) \cdot \Phi(r^\ell) \cdot \mathbf{p}_b$$

Because $\Phi(r^\ell) \cdot \mathbf{p}_b = \mathbf{p}_b$, the r.h.s. simplifies to $\mathbf{p}_i$, and so the edge ij is collapsed.

Multiple connected components The proof of the lemma is completed by considering connected components one at a time to remove the assumption that G is connected. $\square$

5. Rigidity

5.1 Crystallographic and colored frameworks

We now return to the setting of crystallographic frameworks, leading to the proof of Theorem 1 in Section 5.3. The overall structure is very similar to [16, Sections 16–18], but we give sufficient detail for completeness.

5.1.1. Crystallographic frameworks We recall the following definition from the introduction: a *crystallographic framework* $(\tilde{G}, \varphi, \tilde{\boldsymbol{\ell}})$ is given by:

- An infinite graph $\tilde{G}$
- A free action φ on $\tilde{G}$ by a crystallographic group Γ with finite quotient
- An assignment of a *length* ℓ_{ij} to each edge $ij \in E(\tilde{G})$

In what follows, Γ will always be one of the groups Γ_2, Γ_3, Γ_4, or Γ_6.

5.1.2. The realization space A *realization* $\tilde{G}(\mathbf{p}, \Phi)$ of a crystallographic framework $(\tilde{G}, \varphi, \tilde{\boldsymbol{\ell}})$ is given by an assignment $\mathbf{p} = \left(\mathbf{p}_i\right)_{i \in \tilde{V}}$ of points to the vertices of $\tilde{G}$ and a representation Φ of $\Gamma \hookrightarrow \mathrm{Euc}(2)$ by Euclidean isometries acting discretely and co-compactly, such that

$$||\mathbf{p}_i - \mathbf{p}_j|| = \tilde{\boldsymbol{\ell}}_{ij} \qquad \text{for all edges } ij \in \tilde{E} \tag{14}$$

$$\Phi(\gamma) \cdot \mathbf{p}_i = \mathbf{p}_{\gamma(i)} \qquad \text{for all group elements } \gamma \in \Gamma \text{ and vertices } i \in \tilde{V} \tag{15}$$

We see that (15) implies that, to be realizable at all, the framework $(\tilde{G}, \varphi, \tilde{\boldsymbol{\ell}})$ must assign the same length to each edge in every Γ-orbit of the action φ. The condition (14) is the standard one from rigidity theory that says the distances between endpoints of each edge realize the specified lengths.

We define the *realization space* $\mathcal{R}(\tilde{G}, \varphi, \tilde{\boldsymbol{\ell}})$ (shortly $\mathcal{R}$) of a crystallographic framework to be the set of all realizations

$$\mathcal{R}(\tilde{G}, \varphi, \tilde{\boldsymbol{\ell}}) = \left\{(\mathbf{p}, \Phi) : \tilde{G}(\mathbf{p}, \Phi) \text{ is a realization of } (\tilde{G}, \varphi, \tilde{\boldsymbol{\ell}})\right\}$$

5.1.3. The configuration space The group $\mathrm{Euc}(2)$ of Euclidean isometries acts naturally on the realization space. Let $\psi \in \mathrm{Euc}(2)$ be an isometry. For any point $(\mathbf{p}, \Phi) \in \mathcal{R}$,

$$(\psi \circ \mathbf{p}, \Phi^{\psi})$$

is a point in $\mathcal{R}$ as well where Φ^{ψ} is the representation defined by

$$\Phi^{\psi}(\gamma) = \psi \Phi(\gamma) \psi^{-1}.$$

We define the *configuration space* $\mathcal{C}(\tilde{G}, \varphi, \tilde{\boldsymbol{\ell}})$ (shortly $\mathcal{C}$) of a crystallographic framework to be the quotient $\mathcal{R}/\,\mathrm{Euc}(2)$ of the realization space by Euclidean isometries.

Since the spaces $\mathcal{R}$ and $\mathcal{C}$ are subsets of an infinite-dimensional space, there are some technical details to check that we omit in the interest of brevity. Interested readers can find a development for the periodic setting in [12, Appendix A][2]. The present crystallographic case proceeds along the same lines.

5.1.4. Rigidity and flexibility A realization $\tilde{G}(\mathbf{p}, \Phi)$ is defined to be (continuously) *rigid* if it is isolated in the configuration space $\mathcal{C}$. Otherwise it is *flexible*. As the definition makes clear, rigidity is a *local* property that depends on a realization. A framework that is rigid, but ceases to be so if any orbit of bars is removes is defined to be *minimally rigid*.

5.1.5. Colored crystallographic frameworks In principle, the realization and configuration spaces $\mathcal{R}(\tilde{G}, \varphi, \tilde{\boldsymbol{\ell}})$ and $\mathcal{C}(\tilde{G}, \varphi, \tilde{\boldsymbol{\ell}})$ of crystallographic frameworks could be complicated infinite dimensional objects. They are, in fact, equivalent to the finite-dimensional configuration spaces of *colored crystallographic frameworks*, which will be technically simpler to work with. (See Proposition 5.2 below.)

A *colored crystallographic framework* (shortly a *colored framework*) is a triple $(G, \gamma, \boldsymbol{\ell})$, where (G, γ) is a Γ_k-colored graph and $\boldsymbol{\ell} = (\ell_{ij})_{ij \in E(G)}$ is an assignment of a length to each edge. There is a dictionary between crystallographic and colored frameworks, which is a simple modification of the dictionary for direction networks.

5.1.6. The colored realization and configuration spaces A *realization* $G(\mathbf{p}, \Phi)$ of a colored framework is an assignment of points $\mathbf{p} = (\mathbf{p}_i)_{i \in V(G)}$ and a representation Φ of Γ_k by Euclidean isometries acting discretely and cocompactly such that

$$||\Phi(\gamma_{ij}) \cdot \mathbf{p}_j - \mathbf{p}_i||^2 = \ell_{ij}^2$$

for all edges $ij \in E(G)$. The *realization space* $\mathcal{R}(G, \gamma, \ell)$ is then defined to be

$$\mathcal{R}(G, \gamma, \ell) = \{(\mathbf{p}, \Phi) : G(\mathbf{p}, \Phi) \text{ is a realization of } (G, \gamma, \boldsymbol{\ell})\}$$

The Euclidean group $\mathrm{Euc}(2)$ acts naturally on $\mathcal{R}(G, \gamma, \ell)$ by

$$\psi \cdot (\mathbf{p}, \Phi) = (\psi \cdot \mathbf{p}, \Phi^{\psi})$$

where ψ is a Euclidean isometry. Thus we define the *configuration space* $\mathcal{C}(G, \gamma, \ell)$ to be the quotient $\mathcal{R}(G, \gamma, \ell)/\mathrm{Euc}(2)$ of the realization space by the Euclidean group.

5.1.7. The modified configuration space Because it is technically simpler, we will consider the modified realization space $\mathcal{R}'(G, \gamma, \ell)$, which we define to be:

$$\mathcal{R}'(G, \gamma, \ell) = \{(\mathbf{p}, \Phi) : G(\mathbf{p}, \Phi) \text{ is a realization of } (G, \gamma, \boldsymbol{\ell}) \text{ with } \Phi(r_k) \text{ fixing the origin}\}$$

Recall that r_k is the rotation of order k that is one of the generators of Γ_k. The modified configuration space $\mathcal{C}'(G, \gamma, \ell)$ is then defined to be the quotient $\mathcal{R}'(G, \gamma, \ell)/O(2)$ of the modified realization space by the orthogonal group $O(2)$. Since every representation $\Phi \in \mathrm{Rep}(\Gamma_k)$ is conjugate by a Euclidean translation to a representation Φ' that has the origin as a rotation center, this next lemma follows immediately.

[2] The reference [12] is an earlier version of [16].

Lemma 5.1. *Let $(G, \gamma, \boldsymbol{\ell})$ be a colored framework. Then the configuration space $\mathcal{C}(G, \gamma, \ell)$ is homeomorphic to the modified configuration space $\mathcal{C}'(G, \gamma, \ell)$.*

From the definition and Lemma 2.3 we see that the modified configuration space is an algebraic subset of $\mathbb{R}^{2n} \times \mathbb{R}^4$, for Γ_2 and of $\mathbb{R}^{2n} \times \mathbb{R}^2$ for Γ_k with $k = 3, 4, 6$.

5.1.8. Colored rigidity and flexibility We now can define rigidity and flexibility in terms of colored frameworks. A realization $G(\mathbf{p}, \Phi)$ of a colored framework is *rigid* if it is isolated in the configuration space and otherwise *flexible*. Lemma 5.1 implies that a realization is rigid if and only if it is isolated in the modified configuration space.

5.1.9. Equivalence of crystallographic and colored rigidity The connection between the rigidity of crystallographic and colored frameworks is captured in the following proposition, which says that we can switch between the two models.

Proposition 5.2. *Let $(\tilde{G}, \varphi, \tilde{\boldsymbol{\ell}})$ be a crystallographic framework and let $(G, \gamma, \boldsymbol{\ell})$ be an associated colored framework quotient. Then the configuration spaces $\mathcal{C}(\tilde{G}, \varphi, \tilde{\boldsymbol{\ell}})$ and $\mathcal{C}'(G, \gamma, \ell)$ are homeomorphic.*

Proof. This follows from the definitions, a straightforward computation, and Lemma 5.1. □

5.2 Infinitesimal and generic rigidity

As discussed above, the modified realization space $\mathcal{R}'(G, \gamma, \boldsymbol{\ell})$ of a colored framework is an algebraic subset of $\mathbb{R}^{2n+2r}$, where $r = \mathrm{rep}_{\Gamma_k}(\Gamma_k)$. The coordinates are given as follows:

- The first $2n$ coordinates are the coordinates of the points $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_n$
- The final $2r$ coordinates are the vectors v_i specifying the representation of the translation subgroup $\Lambda(\Gamma_k)$. (Since we have "pinned" a rotation center to the origin, the vector w from Lemma 2.3 is fixed to be 0.)

5.2.1. Infinitesimal rigidity As is typical in the derivation of Laman-type theorems, we linearize the problem by considering the tangent space of $\mathcal{R}'(G, \gamma, \boldsymbol{\ell})$ near a realization $G(\mathbf{p}, \Phi)$.

The vectors in the tangent space are infinitesimal motions of the framework, and they can be characterized as follows. Let $(\mathbf{q}, \mathbf{u}_1, \mathbf{u}_2) \in \mathbb{R}^{2n+4}$ for $k = 2$ or $(\mathbf{q}, \mathbf{u}_1) \in \mathbb{R}^{2n+2}$ for $k = 3, 4, 6$. To this vector there is an associated representation Φ' defined by $\Phi'(r_k) = (0, R_k)$ and $\Phi'(t_i) = (\mathbf{u}_i, \mathrm{Id})$. Then differentiation of the length equations yield this linear system ranging over all edges $ij \in E(G)$:

$$\left\langle \Phi(\gamma_{ij}) \cdot \mathbf{p}_j - \mathbf{p}_i, \Phi'(\gamma_{ij}) \cdot \mathbf{q}_j - \mathbf{q}_i \right\rangle \tag{16}$$

The given data are the $\mathbf{p}_i$ and Φ, and then unknowns are the $\mathbf{q}_i$ and Φ'. A realization $G(\mathbf{p}, \Phi)$ of a colored framework is defined to be *infinitesimally rigid* if the system (16) has a 1-dimensional solution space. A realization that is infinitesimally rigid but ceases to be so when any colored edge is removed is minimally infinitesimally rigid.

5.2.2. Infinitesimal rigidity implies rigidity A standard kind of result relating infinitesimal rigidity and rigidity for generic frameworks holds in our setting. Since our realization space is finite dimensional, we can adapt the arguments of, e.g., [1] to our setting to show:

Lemma 5.3. *If a realization $G(\mathbf{p},\Phi)$ of a colored framework is infinitesimally rigid, then it is rigid.*

5.2.3. Generic rigidity The converse of Lemma 5.3 does not hold in general, but it does for nearly all realizations. Let $(G,\gamma,\boldsymbol{\ell})$ be a colored framework. A realization $G(\mathbf{p},\Phi)$ is defined to be *regular* for $(G,\gamma,\boldsymbol{\ell})$ if the rank of the system (16) is maximal over all realizations.

Whether a realization is regular depends on both the colored graph (G,γ) and the given lengths $\boldsymbol{\ell}$. Let $G(\mathbf{p},\Phi)$ be a regular realization of a colored framework. If, in addition, the rank of (16) at $G(\mathbf{p},\Phi)$ is maximal over all realizations of colored frameworks with the same colored graph (G,γ), we define $G(\mathbf{p},\Phi)$ to be *generic*. We define the rank of (16) at a generic realization to be its *generic rank*. Since it depends on formal minors of the matrix underlying (16) only, it is a property of the colored graph (G,γ).

If $(G,\gamma,\boldsymbol{\ell})$ is a framework with generic realizations, it is immediate that the set of non-generic realizations is a proper algebraic subset of the realization space. Alternatively, if we consider frameworks as being induced by realizations, the set of non-generic realizations is a proper algebraic subset of $\mathbb{R}^{2n+2r}$, where $r=1$ for Γ_3, Γ_4, and Γ_6, and $r=2$ for Γ_2.

For generic realizations, a standard argument (again, along the lines of [1]) shows that rigidity and infinitesimal rigidity coincide.

Proposition 5.4. *A generic realization of a colored framework $(G,\gamma,\boldsymbol{\ell})$ is rigid if and only if it is infinitesimally rigid.*

5.3 Proof of Theorem 1

We recall, from the introduction, our main theorem:

Theorem 1. *Let Γ be an orientation-preserving crystallographic group. A generic crystallographic framework $(\tilde{G},\varphi,\tilde{\ell})$ with symmetry group Γ is minimally rigid if and only if its colored quotient graph is Γ-colored-Laman.*

The proof occupies the rest of this section.

5.3.1. Reduction to colored frameworks By Proposition 5.2, it is sufficient to prove the statement of Theorem 1 for colored frameworks. Proposition 5.4 then implies that the theorem will follow from a characterization of generic infinitesimal rigidity for colored frameworks.

Thus, to prove the theorem, we show that, for a colored graph (G,γ) with n vertices and $m = 2n + \mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma_k)) - 1$ edges, the generic rank of the system (16) is m if and only if (G,γ) is a Γ-colored-Laman graph.

5.3.2. Necessity: the "Maxwell direction" We recall the definition of the sparsity function $h(G)$ from Section 3.4, which defines Γ-colored-Laman graphs. We have, for a colored graph (G,γ) with n vertices and c connected components $G_1, G_2, \ldots, G_c$,

$$h(G) = 2n + \mathrm{rep}_{\Gamma_k}(G) - 1 - \sum_{i=1}^{c} T(G_i)$$

Proposition 5.5. *Let (G,γ) be a colored graph. Then the generic rank of the system* (16) *is at most* $h(G)$.

Proof. Let $G(\mathbf{p},\Phi)$ be any realization of a colored framework on a colored graph (G,γ) with no collapsed edges. That is select a representation Φ of Γ_k and points $\mathbf{p}_i$, such that, $\Phi(\gamma_{ij})\cdot\mathbf{p}_j \neq \mathbf{p}_i$ for all edges $ij \in E(G)$.

We now define the direction $\mathbf{d}_{ij}$ to be $(\Phi(\gamma_{ij})\cdot\mathbf{p}_j - \mathbf{p}_i)^\perp$ for each edge $ij \in E(G)$. These directions define a colored direction network $(G,\gamma,\mathbf{d})$ with the property that any solution to this direction network corresponds to an infinitesimal motion of the colored framework realization $G(\mathbf{p},\Phi)$.

Lemma 4.13 implies that there are

$$\mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma_k)) - \mathrm{rep}_{\Gamma_k}(G) + \sum_{i=1}^{c} T(G_i)$$

dimensions of realizations with every edge collapsed. By construction, there is a non-collapsed realization of this direction network as well: it is simply $(\mathbf{p},\Phi)$ rotated by $\pi/2$. Since this is not obtained by taking linear combinations of realizations where every edge is collapsed, the dimension of the space of infinitesimal motions is always at least

$$\mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma_k)) - \mathrm{rep}_{\Gamma_k}(G) + \sum_{i=1}^{c} T(G_i) + 1$$

The proposition follows by subtracting from $2n + \mathrm{rep}_{\Gamma_k}(\Lambda(\Gamma_k))$ and comparing to $h(G)$. $\square$

5.3.3. Sufficiency: the "Laman direction" The other direction of the proof of Theorem 1 is this next proposition

Proposition 5.6. *Let (G,γ) be a Γ-colored-Laman graph. Then the generic rank of the system* (16) *is* $h(G)$.

Proof. It is sufficient to construct a single example at which this rank is attained, since the generic rank is always at least the rank for any specific realization. We will do this using direction networks.

Let (G,γ) be a Γ-colored-Laman graph, and select a direction $\mathbf{d}_{ij}$ for each edge $ij \in E(G)$, such that both $\mathbf{d}$ and $\mathbf{d}^\perp = (\mathbf{d}_{ij}^\perp)$ are generic in the sense of Proposition 4.1. By Proposition 4.1, the colored direction network $(G,\gamma,\mathbf{d})$ has a unique, up to scaling, faithful realization $(\mathbf{p},\Phi)$, which implies that, for all edges $ij \in E(G)$

$$\Phi(\gamma_{ij})\cdot\mathbf{p}_j - \mathbf{p}_i = \alpha_{ij}\mathbf{d}_{ij}$$

for some non-zero scalar $\alpha_{ij} \in \mathbb{R}$. It follows that, by replacing $\mathbf{d}_{ij}$ with $\Phi(\gamma_{ij})\cdot\mathbf{p}_j - \mathbf{p}_i$ in the direction realization system (9) we obtain (16). Since $\mathbf{d}^\perp$ is also generic for Proposition 4.1, we conclude that (16) has full rank as desired. $\square$

References

[1] L. Asimow and B. Roth. The rigidity of graphs. *Trans. Amer. Math. Soc.*, 245:279–289, 1978. ISSN 0002-9947. doi: 10.2307/1998867. URL http://dx.doi.org/10.2307/1998867.

[2] L. Bieberbach. Über die Bewegungsgruppen der Euklidischen Räume. *Math. Ann.*, 70(3): 297–336, 1911. doi: 10.1007/BF01564500.

[3] L. Bieberbach. Über die Bewegungsgruppen der Euklidischen Räume (Zweite Abhandlung.) Die Gruppen mit einem endlichen Fundamentalbereich. *Math. Ann.*, 72(3):400–412, 1912. doi: 10.1007/BF01456724.

[4] C. S. Borcea and I. Streinu. Periodic frameworks and flexibility. *Proc. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.*, 466(2121):2633–2649, 2010. ISSN 1364-5021. doi: 10.1098/rspa.2009.0676. URL http://dx.doi.org/10.1098/rspa.2009.0676.

[5] C. S. Borcea and I. Streinu. Minimally rigid periodic graphs. *Bulletin of the London Mathematical Society*, 2011. doi: 10.1112/blms/bdr044. URL http://dx.doi.org/blms.bdr044.abstract.

[6] J. H. Conway, O. Delgado Friedrichs, D. H. Huson, and W. P. Thurston. On three-dimensional space groups. *Beiträge Algebra Geom.*, 42(2):475–507, 2001. ISSN 0138-4821.

[7] J. Edmonds. Minimum partition of a matroid into independent subsets. *J. Res. Nat. Bur. Standards Sect. B*, 69B:67–72, 1965. ISSN 0160-1741.

[8] J. Edmonds and G.-C. Rota. Submodular set functions (abstract). In *Waterloo Combinatorics Conference*, University of Waterloo, Ontario, 1966.

[9] A. Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002. ISBN 0-521-79160-X; 0-521-79540-0.

[10] A. Lee and I. Streinu. Pebble game algorithms and sparse graphs. *Discrete Math.*, 308 (8):1425–1437, 2008. ISSN 0012-365X. doi: 10.1016/j.disc.2007.07.104. URL http://dx.doi.org/10.1016/j.disc.2007.07.104.

[11] L. Lovász and Y. Yemini. On generic rigidity in the plane. *SIAM J. Algebraic Discrete Methods*, 3(1):91–98, 1982. ISSN 0196-5212. doi: 10.1137/0603009. URL http://dx.doi.org/10.1137/0603009.

[12] J. Malestein and L. Theran. Generic combinatorial rigidity of periodic frameworks. Preprint, arXiv:1008.1837v2, 2010. URL http://arxiv.org/abs/1008.1837v2.

[13] J. Malestein and L. Theran. Generic rigidity of frameworks with orientation-preserving crystallographic symmetry. Preprint, arXiv:1108.2518, 2011. URL http://arxiv.org/abs/1108.2518.

[14] J. Malestein and L. Theran. Generic rigidity of reflection frameworks. Preprint, arXiv:1203.2276, 2012. URL http://arxiv.org/abs/1203.2276.

[15] J. Malestein and L. Theran. Frameworks with forced symmetry I: rotations and reflections. Submitted manuscript, 2012. See also [13].

[16] J. Malestein and L. Theran. Generic combinatorial rigidity of periodic frameworks. *Advances in Mathematics*, 233:291–331, 2013. doi: 10.1016/j.aim.2012.10.007. URL `http://dx.doi.org/10.1016/j.aim.2012.10.007`.

[17] J. Oxley. *Matroid theory*, volume 21 of *Oxford Graduate Texts in Mathematics*. Oxford University Press, Oxford, second edition, 2011. ISBN 978-0-19-960339-8.

[18] A. Recski. A network theory approach to the rigidity of skeletal structures. II. Laman's theorem and topological formulae. *Discrete Appl. Math.*, 8(1):63–68, 1984. ISSN 0166-218X. doi: 10.1016/0166-218X(84)90079-9. URL `http://dx.doi.org/10.1016/0166-218X(84)90079-9`.

[19] E. Ross. *The Rigidity of Periodic Frameworks as Graphs on a Torus*. PhD thesis, York University, 2011. URL `http://www.math.yorku.ca/~ejross/RossThesis.pdf`.

[20] E. Ross, B. Schulze, and W. Whiteley. Finite motions from periodic frameworks with added symmetry. *International Journal of Solids and Structures*, 48(11-12):1711 – 1729, 2011. ISSN 0020-7683. doi: 10.1016/j.ijsolstr.2011.02.018. URL `http://www.sciencedirect.com/science/article/pii/S0020768311000850`.

[21] B. Schulze. Symmetric versions of Laman's theorem. *Discrete Comput. Geom.*, 44(4):946–972, 2010. ISSN 0179-5376. doi: 10.1007/s00454-009-9231-x. URL `http://dx.doi.org/10.1007/s00454-009-9231-x`.

[22] B. Schulze. Symmetric Laman theorems for the groups $\mathcal{C}_2$ and $\mathcal{C}_s$. *Electron. J. Combin.*, 17 (1):Research Paper 154, 61, 2010. ISSN 1077-8926. URL `http://www.combinatorics.org/Volume_17/Abstracts/v17i1r154.html`.

[23] B. Schulze and W. Whiteley. The orbit rigidity matrix of a symmetric framework. *Discrete & Computational Geometry*, 46:561–598, 2011. ISSN 0179-5376. doi: 10.1007/s00454-010-9317-5. URL `http://dx.doi.org/10.1007/s00454-010-9317-5`.

[24] I. Streinu and L. Theran. Slider-pinning rigidity: a Maxwell-Laman-type theorem. *Discrete & Computational Geometry*, 44(4):812–837, 2010. ISSN 0179-5376. doi: 10.1007/s00454-010-9283-y. URL `http://dx.doi.org/10.1007/s00454-010-9283-y`.

[25] S.-I. Tanigawa. Matroids of gain graphs in applied discrete geometry. Preprint, arXiv:1207.3601, 2012. URL `http://arxiv.org/abs/1207.3601`.

[26] W. Whiteley. The union of matroids and the rigidity of frameworks. *SIAM J. Discrete Math.*, 1(2):237–255, 1988. ISSN 0895-4801. doi: 10.1137/0401025. URL `http://dx.doi.org/10.1137/0401025`.

[27] T. Zaslavsky. Voltage-graphic matroids. In *Matroid theory and its applications*, pages 417–424. Liguori, Naples, 1982.

[28] T. Zaslavsky. A mathematical bibliography of signed and gain graphs and allied areas. *Electron. J. Combin.*, 5:Dynamic Surveys 8, 124 pp. (electronic), 1998. ISSN 1077-8926. URL `http://www.combinatorics.org/Surveys/index.html`. Manuscript prepared with Marge Pratt.